



(12) **United States Patent**
Anderson et al.

(10) **Patent No.:** **US 11,086,784 B2**
(45) **Date of Patent:** ***Aug. 10, 2021**

(54) **INVALIDATING TRACK FORMAT INFORMATION FOR TRACKS IN CACHE**

(58) **Field of Classification Search**
CPC G06F 3/0689; G06F 12/0246; G06F 12/0833; G06F 12/123

(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION**, Armonk, NY (US)

See application file for complete search history.

(56) **References Cited**

(72) Inventors: **Kyler A. Anderson**, Sahuarita, AZ (US); **Kevin J. Ash**, Tucson, AZ (US); **Lokesh M. Gupta**, Tucson, AZ (US); **Matthew J. Kalos**, Tucson, AZ (US)

U.S. PATENT DOCUMENTS
5,636,359 A 6/1997 Beardsley et al.
6,314,490 B1 ‡ 11/2001 Morein G06F 12/0864 711/12

(Continued)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

FOREIGN PATENT DOCUMENTS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

CN 106294772 A ‡ 1/2017
CN 106294772 A 1/2017

This patent is subject to a terminal disclaimer.

OTHER PUBLICATIONS

U.S. Appl. No. 15/672,460, filed Aug. 9, 2017, (18.735).‡

(Continued)

(21) Appl. No.: **16/742,832**

Primary Examiner — Michelle T Bechtold

(22) Filed: **Jan. 14, 2020**

(74) *Attorney, Agent, or Firm* — Konrad Raynes Davda & Victor LLP; David W. Victor

(65) **Prior Publication Data**
US 2020/0151105 A1 May 14, 2020

(57) **ABSTRACT**

Related U.S. Application Data

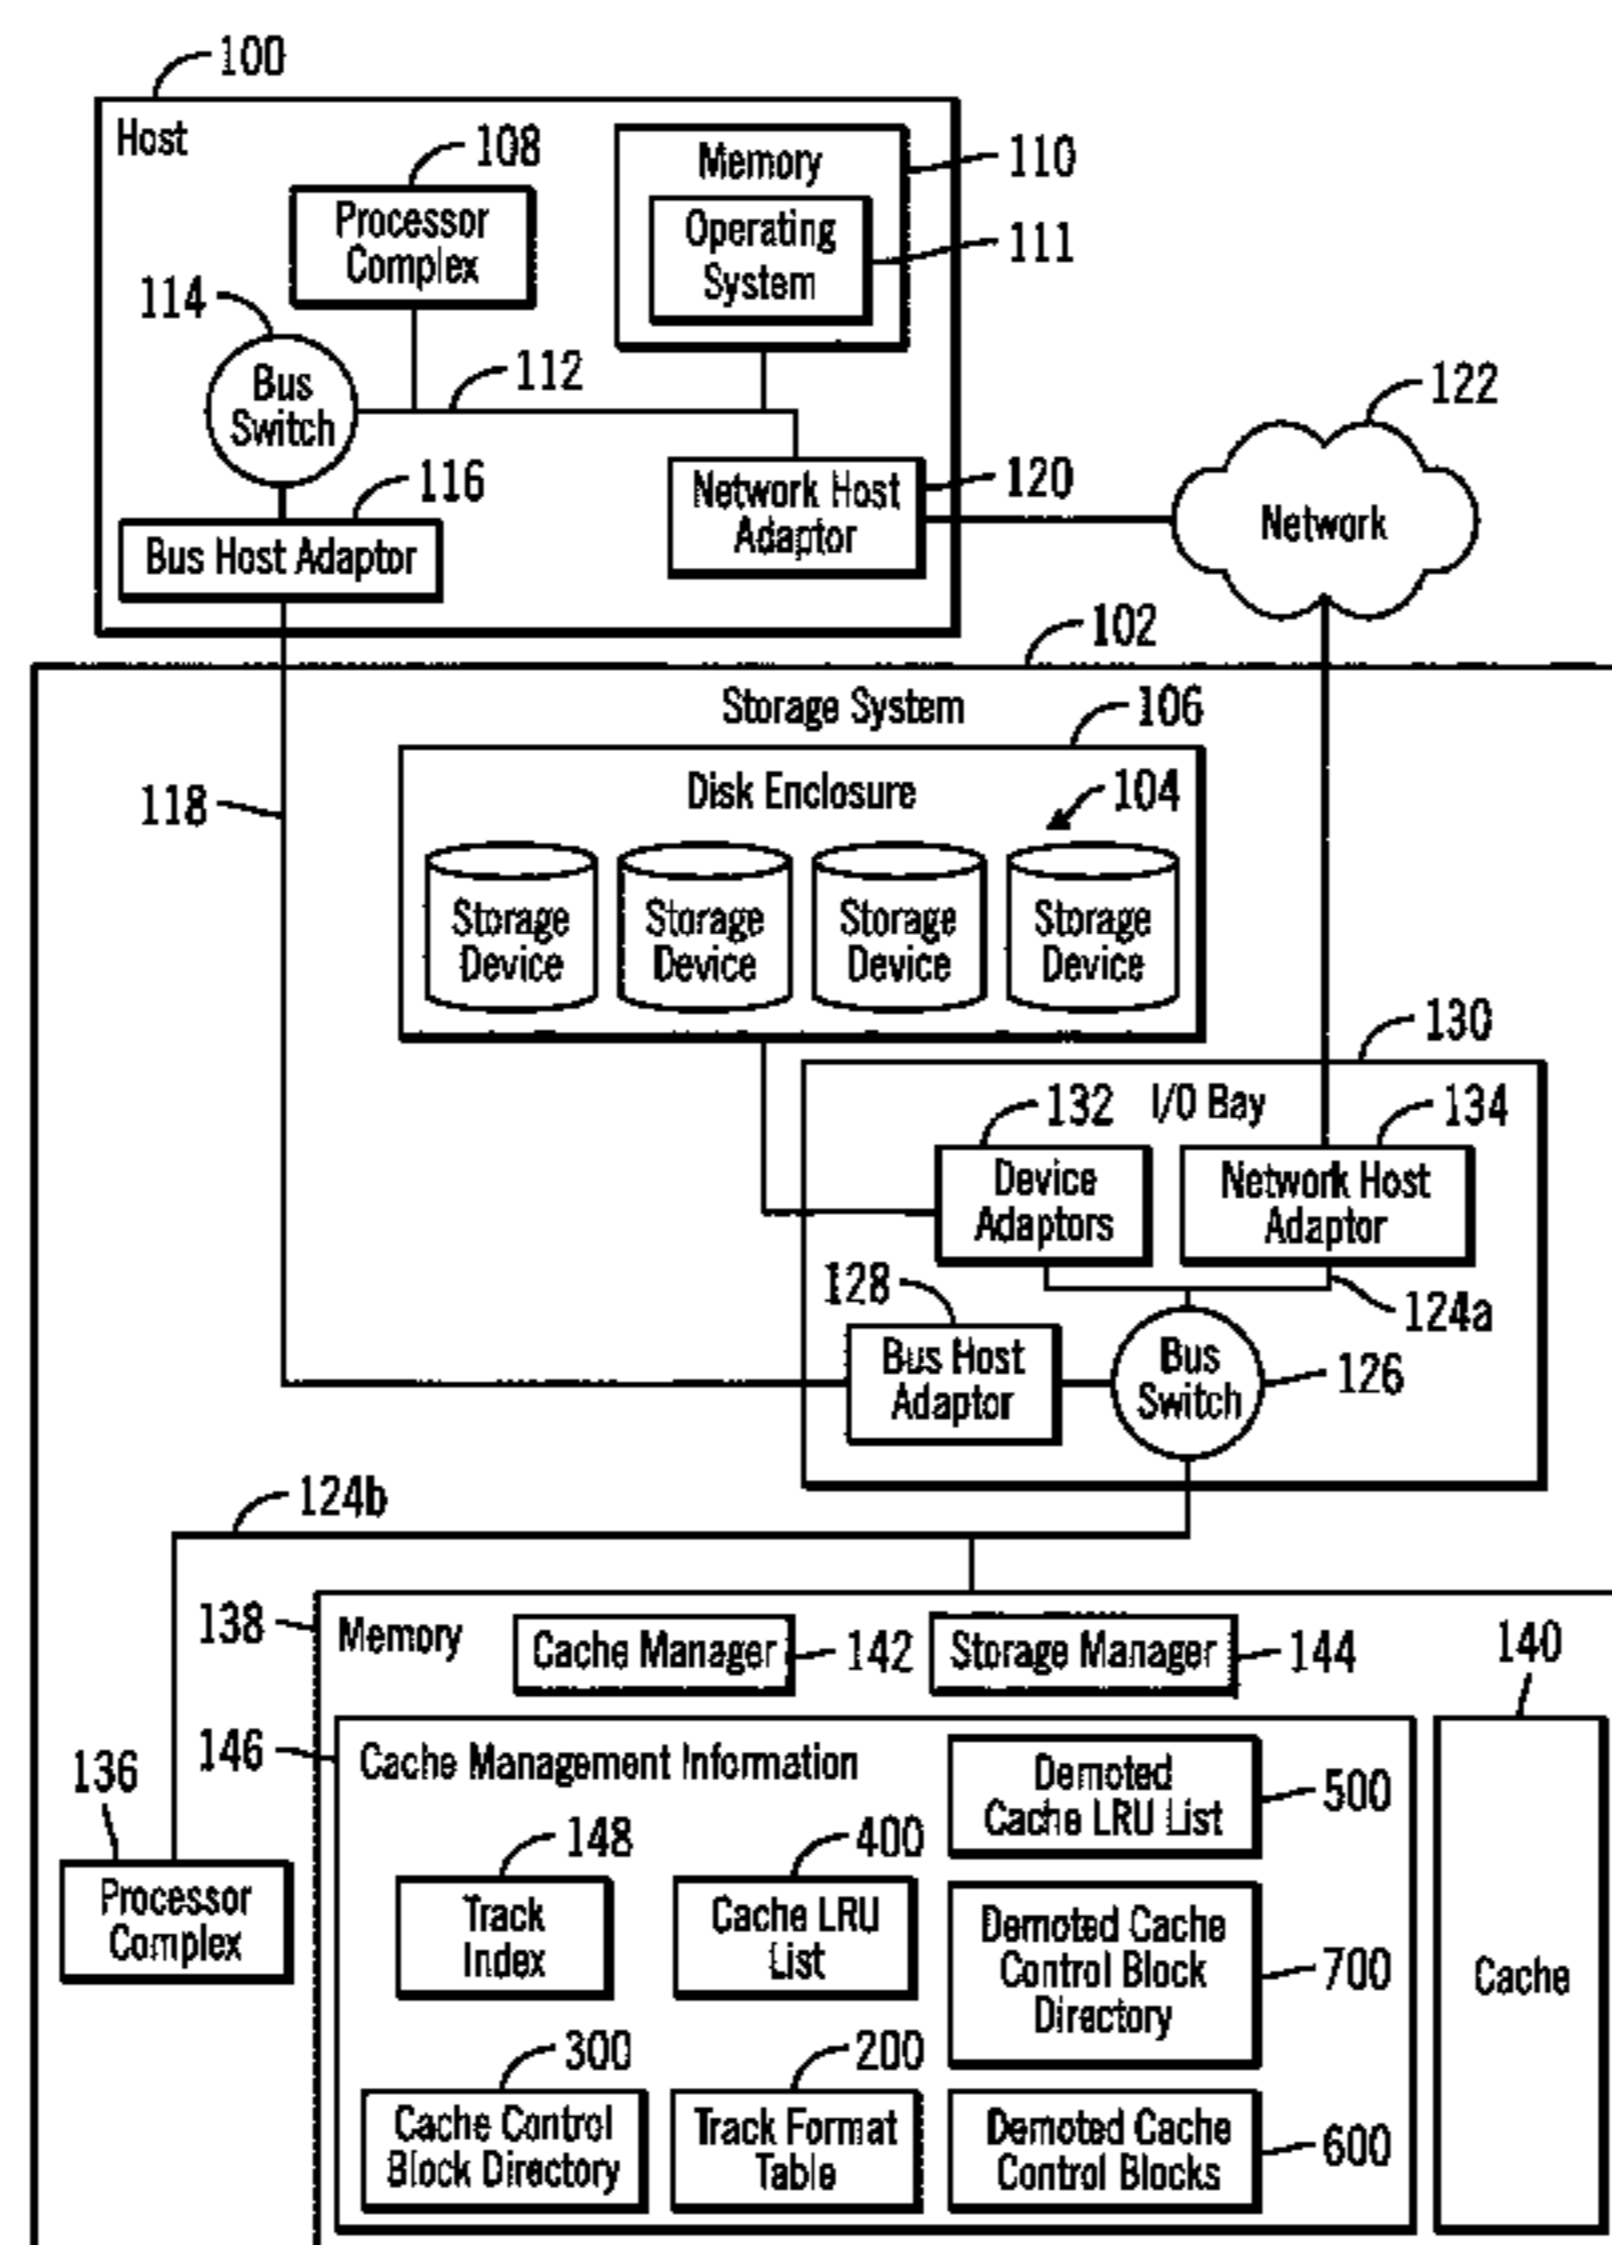
(63) Continuation of application No. 15/672,470, filed on Aug. 9, 2017, now Pat. No. 10,579,532.

Provided are a computer program product, system, and method for invalidating track format information for tracks in cache. Demoted tracks demoted from the cache are indicated in a demoted track list. Track format information is saved for the demoted tracks. The track format information indicates a layout of data in the demoted tracks, wherein the track format information for the demoted tracks is used when the demoted tracks are staged back into the cache. An operation is initiated to invalidate a metadata track of the metadata tracks in the storage. Demoted tracks indicated in the demoted track list having metadata in the metadata track to invalidate are removed. The track format information for the demoted tracks having metadata in the metadata track to invalidate is removed.

(51) **Int. Cl.**
G06F 12/0871 (2016.01)
G06F 12/0831 (2016.01)
(Continued)

(52) **U.S. Cl.**
CPC **G06F 12/0871** (2013.01); **G06F 3/0689** (2013.01); **G06F 12/0246** (2013.01);
(Continued)

21 Claims, 12 Drawing Sheets



- (51) **Int. Cl.**
G06F 12/123 (2016.01)
G06F 12/02 (2006.01)
G06F 3/06 (2006.01)
- (52) **U.S. Cl.**
 CPC **G06F 12/0833** (2013.01); **G06F 12/123**
 (2013.01); **G06F 2212/261** (2013.01); **G06F**
2212/313 (2013.01); **G06F 2212/462** (2013.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,438,661 B1 † 8/2002 Beardsley G06F 11/1435
 711/112

6,832,290 B2 † 12/2004 Todd G06F 16/21
 711/11

7,139,933 B2 † 11/2006 Hsu G06F 11/2092
 714/6

7,434,002 B1 † 10/2008 Zedlewski G06F 11/3409
 711/13

8,027,952 B2 † 9/2011 Kern G06F 11/2058
 707/61

8,332,581 B2 † 12/2012 Candelaria G06F 16/10
 711/11

8,606,767 B2 † 12/2013 Brown G06F 3/0617
 707/69

8,825,973 B2 † 9/2014 Ash G06F 12/0802
 711/11

8,825,975 B2 † 9/2014 Ash G06F 12/0802
 711/11

9,104,552 B1 † 8/2015 Bolt G06F 12/0246

9,298,567 B2 † 3/2016 Blea G06F 11/2069

9,317,207 B2 † 4/2016 Hu G06F 3/0611

9,489,389 B2 † 11/2016 Patil G06F 16/172

9,582,421 B1 † 2/2017 Agarwala G06F 12/0253

10,223,274 B1 3/2019 Anderson et al.

2003/0131163 A1 † 7/2003 Ishii G06F 3/0605
 710/52

2003/0177130 A1 † 9/2003 Todd G06F 16/21

2004/0250022 A1 † 12/2004 Jarvis G06F 12/0866
 711/118

2004/0267706 A1 † 12/2004 Springer, Sr. G06F 3/061

2006/0080510 A1 † 4/2006 Benhase G06F 12/123
 711/13

2007/0245099 A1 † 10/2007 Gray G06F 9/3004
 711/15

2008/0120482 A1 † 5/2008 Jarvis G06F 11/2069
 711/16

2008/0126913 A1 † 5/2008 Benhanokh G06F 11/1435
 714/76

2008/0256286 A1 † 10/2008 Ash G06F 11/2284
 711/6

2011/0191522 A1 † 8/2011 Condict G06F 12/123
 711/10

2012/0102078 A1 † 4/2012 Flick G06F 16/68
 707/82

2012/0260043 A1 † 10/2012 Benhase G06F 3/0619
 711/133

2013/0185489 A1 † 7/2013 Benhase G06F 12/123
 711/10

2013/0185493 A1 † 7/2013 Benhase G06F 12/128
 711/10

2014/0115261 A1 † 4/2014 Maybee G06F 12/0897
 711/13

2014/0258628 A1 † 9/2014 Shivashankaraiah
 G06F 12/0871
 711/12

2014/0325145 A1 † 10/2014 Sampathkumar ... G06F 12/0868
 711/11

2015/0012690 A1 * 1/2015 Bruce G06F 12/0804
 711/103

2015/0286432 A1 † 10/2015 Dain G06F 3/0688
 711/10

2015/0309892 A1 † 10/2015 Ramasubramaniam
 G06F 11/2005
 714/4

2016/0085460 A1 † 3/2016 Chandrashekar G06F 3/067
 711/13

2016/0140053 A1 † 5/2016 Ash G06F 12/123
 711/136

2016/0232102 A1 † 8/2016 Ash G06F 12/123

2017/0052822 A1 † 2/2017 Ash G06F 9/5016

2017/0111468 A1 † 4/2017 Ash H04L 67/2852

2017/0124001 A1 † 5/2017 Ash G06F 12/123

2019/0034302 A1 1/2019 Anderson et al.

2019/0034303 A1 1/2019 Anderson et al.

2019/0034304 A1 † 1/2019 Anderson G06F 11/2012

2019/0034355 A1 † 1/2019 Anderson G06F 12/123

2019/0042096 A1 2/2019 Anderson et al.

2019/0050339 A1 † 2/2019 Anderson G06F 12/0871

2019/0050340 A1 2/2019 Anderson et al.

2019/0073279 A1 3/2019 Anderson et al.

OTHER PUBLICATIONS

U.S. Appl. No. 15/666,525, filed Aug. 1, 2017, (18.738). †

U.S. Appl. No. 15/662,200, filed Jul. 27, 2017, (18.734). †

U.S. Appl. No. 15/662,205, filed Jul. 27, 2017, (18.752). †

U.S. Appl. No. 15/694,704, filed Sep. 1, 2017, (18.756). †

U.S. Appl. No. 15/688,763, filed Aug. 28, 2017, (18.742). †

Anonymously; “Optimized Deduplication Metadata Online Redistrib-
 tion”; <http://ip.com/IPCOM/000228925D>; Jul. 15, 2013. †

U.S. Appl. No. 15/662,194, filed Jul. 27, 2017, (18.733). †

U.S. Appl. No. 15/662,186, filed Jul. 27, 2017, (18.729). †

Anonymously; “A Systematic Framework for Managing Data Rules
 at Scale: for Big Data, the Cloud, and Beyond”; <http://ip.com/IPCOM/000242035D>; Jun. 15, 2015. †

Anonymously; “Record Level Change Tracking Included within a
 Data Set”; <http://ip.com/IPCOM/000239805D>; Dec. 2, 2014. †

IBM; “System, Method and Service for Automated Capacity Expans-
 ion of a Data Deduplication Repository”; <http://ip.com/IPCOM/000184411D>; Jun. 23, 2009. †

List of IBM Patents or Patent Applications Treated as Related, dated
 Oct. 24, 2017, pp. 2. †

Response dated May 5, 2019, pp. 12, to Office Action dated Feb. 7,
 2019, pp. 24, for U.S. Appl. No. 15/662,200, (18.734). †

Office Action dated Feb. 7, 2019, pp. 24, for U.S. Appl. No.
 15/662,200, (18.734). †

List of IBM Patents and Applications Treated as Related, Jul. 11,
 2019, pp. 2. †

Office Action dated Jun. 20, 2019, pp. 51, for U.S. Appl. No.
 15/672,460, (18.735). †

Office Action dated Mar. 20, 2019, pp. 20, for U.S. Appl. No.
 15/662,194, (18.733). †

Response dated Jun. 22, 2019, pp. 15, to Office Action dated Mar.
 20, 2019, pp. 20, for U.S. Appl. No. 15/662,194, (18.733). †

English translation of Chinese patent CN106294772 dated Jan. 4,
 2017, pp. 10. †

International Search Report and Written Opinion dated Nov. 21,
 2018, for Application PCT IB2018/055378, pp. 9. †

Response dated Sep. 20, 2019, pp. 11, to Office Action dated Jun.
 20, 2019, pp. 51, for U.S. Appl. No. 15/672,460, (18.735). †

Preliminary Amendment dated Jan. 14, 2020, pp. 8, for U.S. Appl.
 No. 16/472,832, (18.754C1). †

List of IBM Patents and Applications Treated as Related, Feb. 11,
 2020, pp. 2. †

U.S. Appl. No. 16/472,832, filed Jan. 14, 2020, (18.754C1). †

Final Office Action dated Dec. 31, 2019, pp. 57, for U.S. Appl. No.
 15/672,460, (18.735). †

Notice of Allowance dated Aug. 21, 2019, pp. 20, for U.S. Appl. No.
 15/662,200, (18.734). †

Notice of Allowance dated Nov. 29, 2019, pp. 13, for U.S. Appl. No.
 15/662,194, (18.733). †

Amendment dated Nov. 8, 2019, pp. 12, to Final Office Action dated
 Sep. 12, 2019, pp. 19, for U.S. Appl. No. 15/662,194, (18.733). †

(56)

References Cited

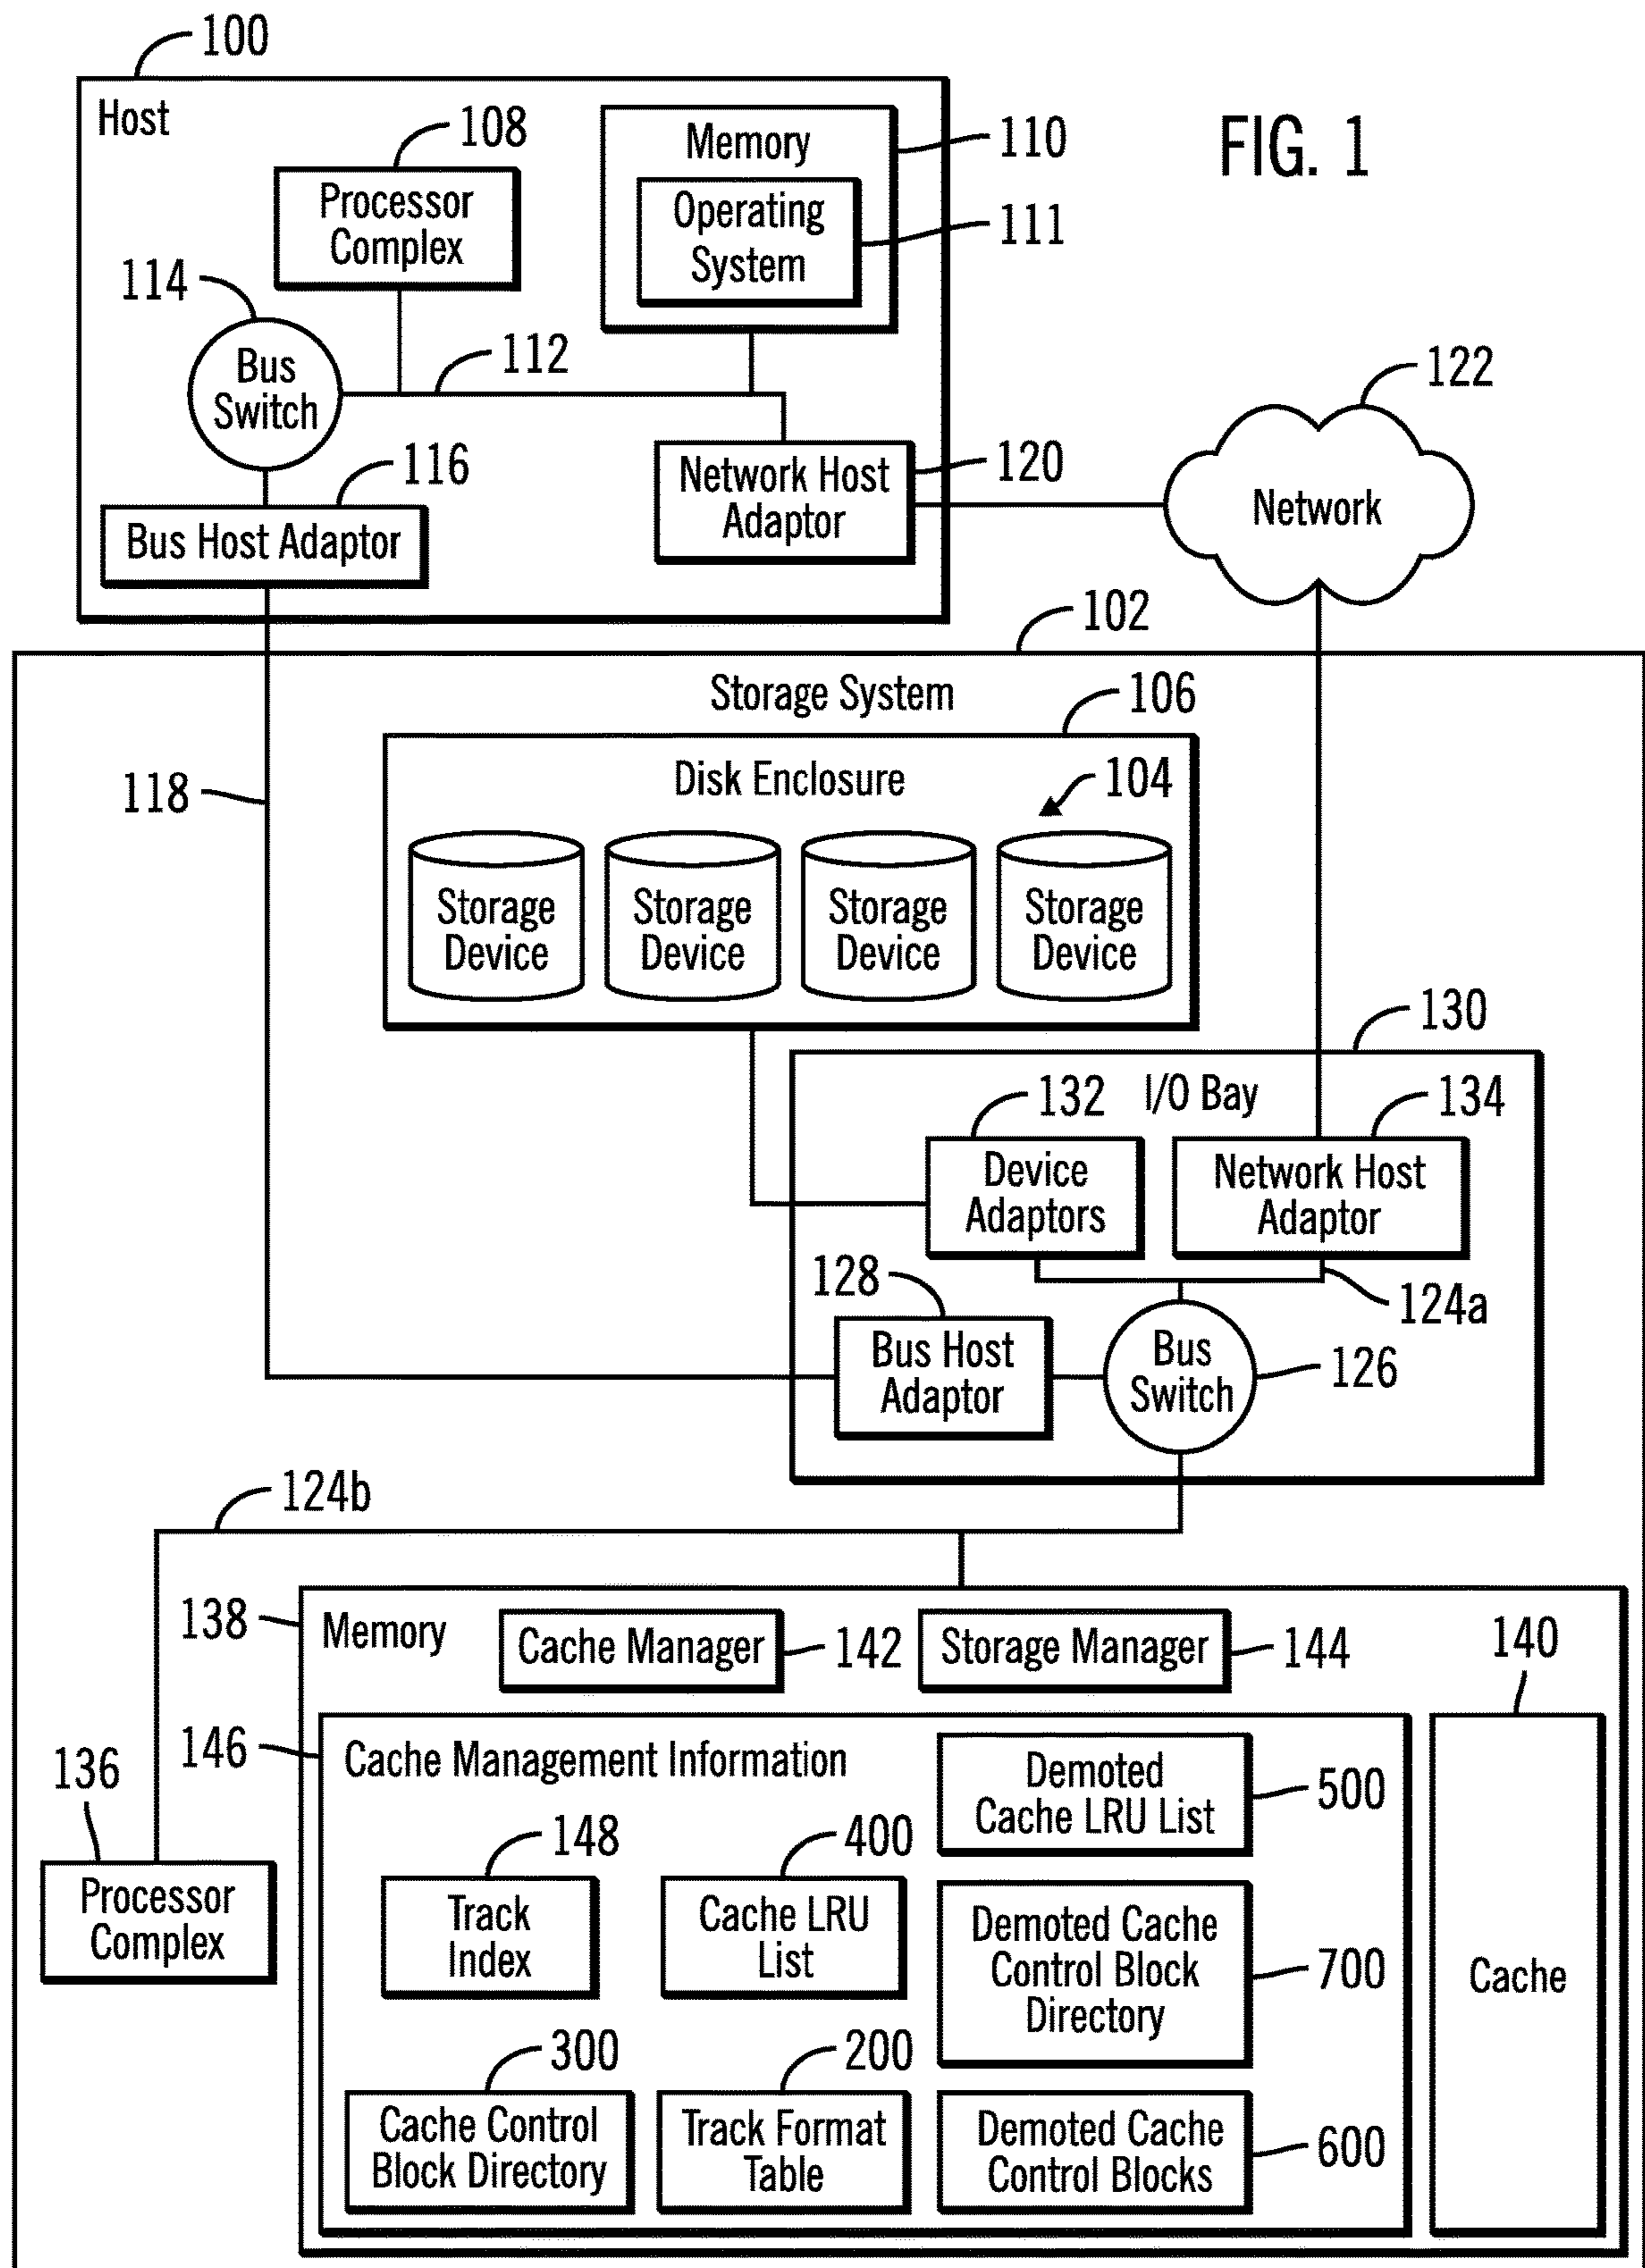
OTHER PUBLICATIONS

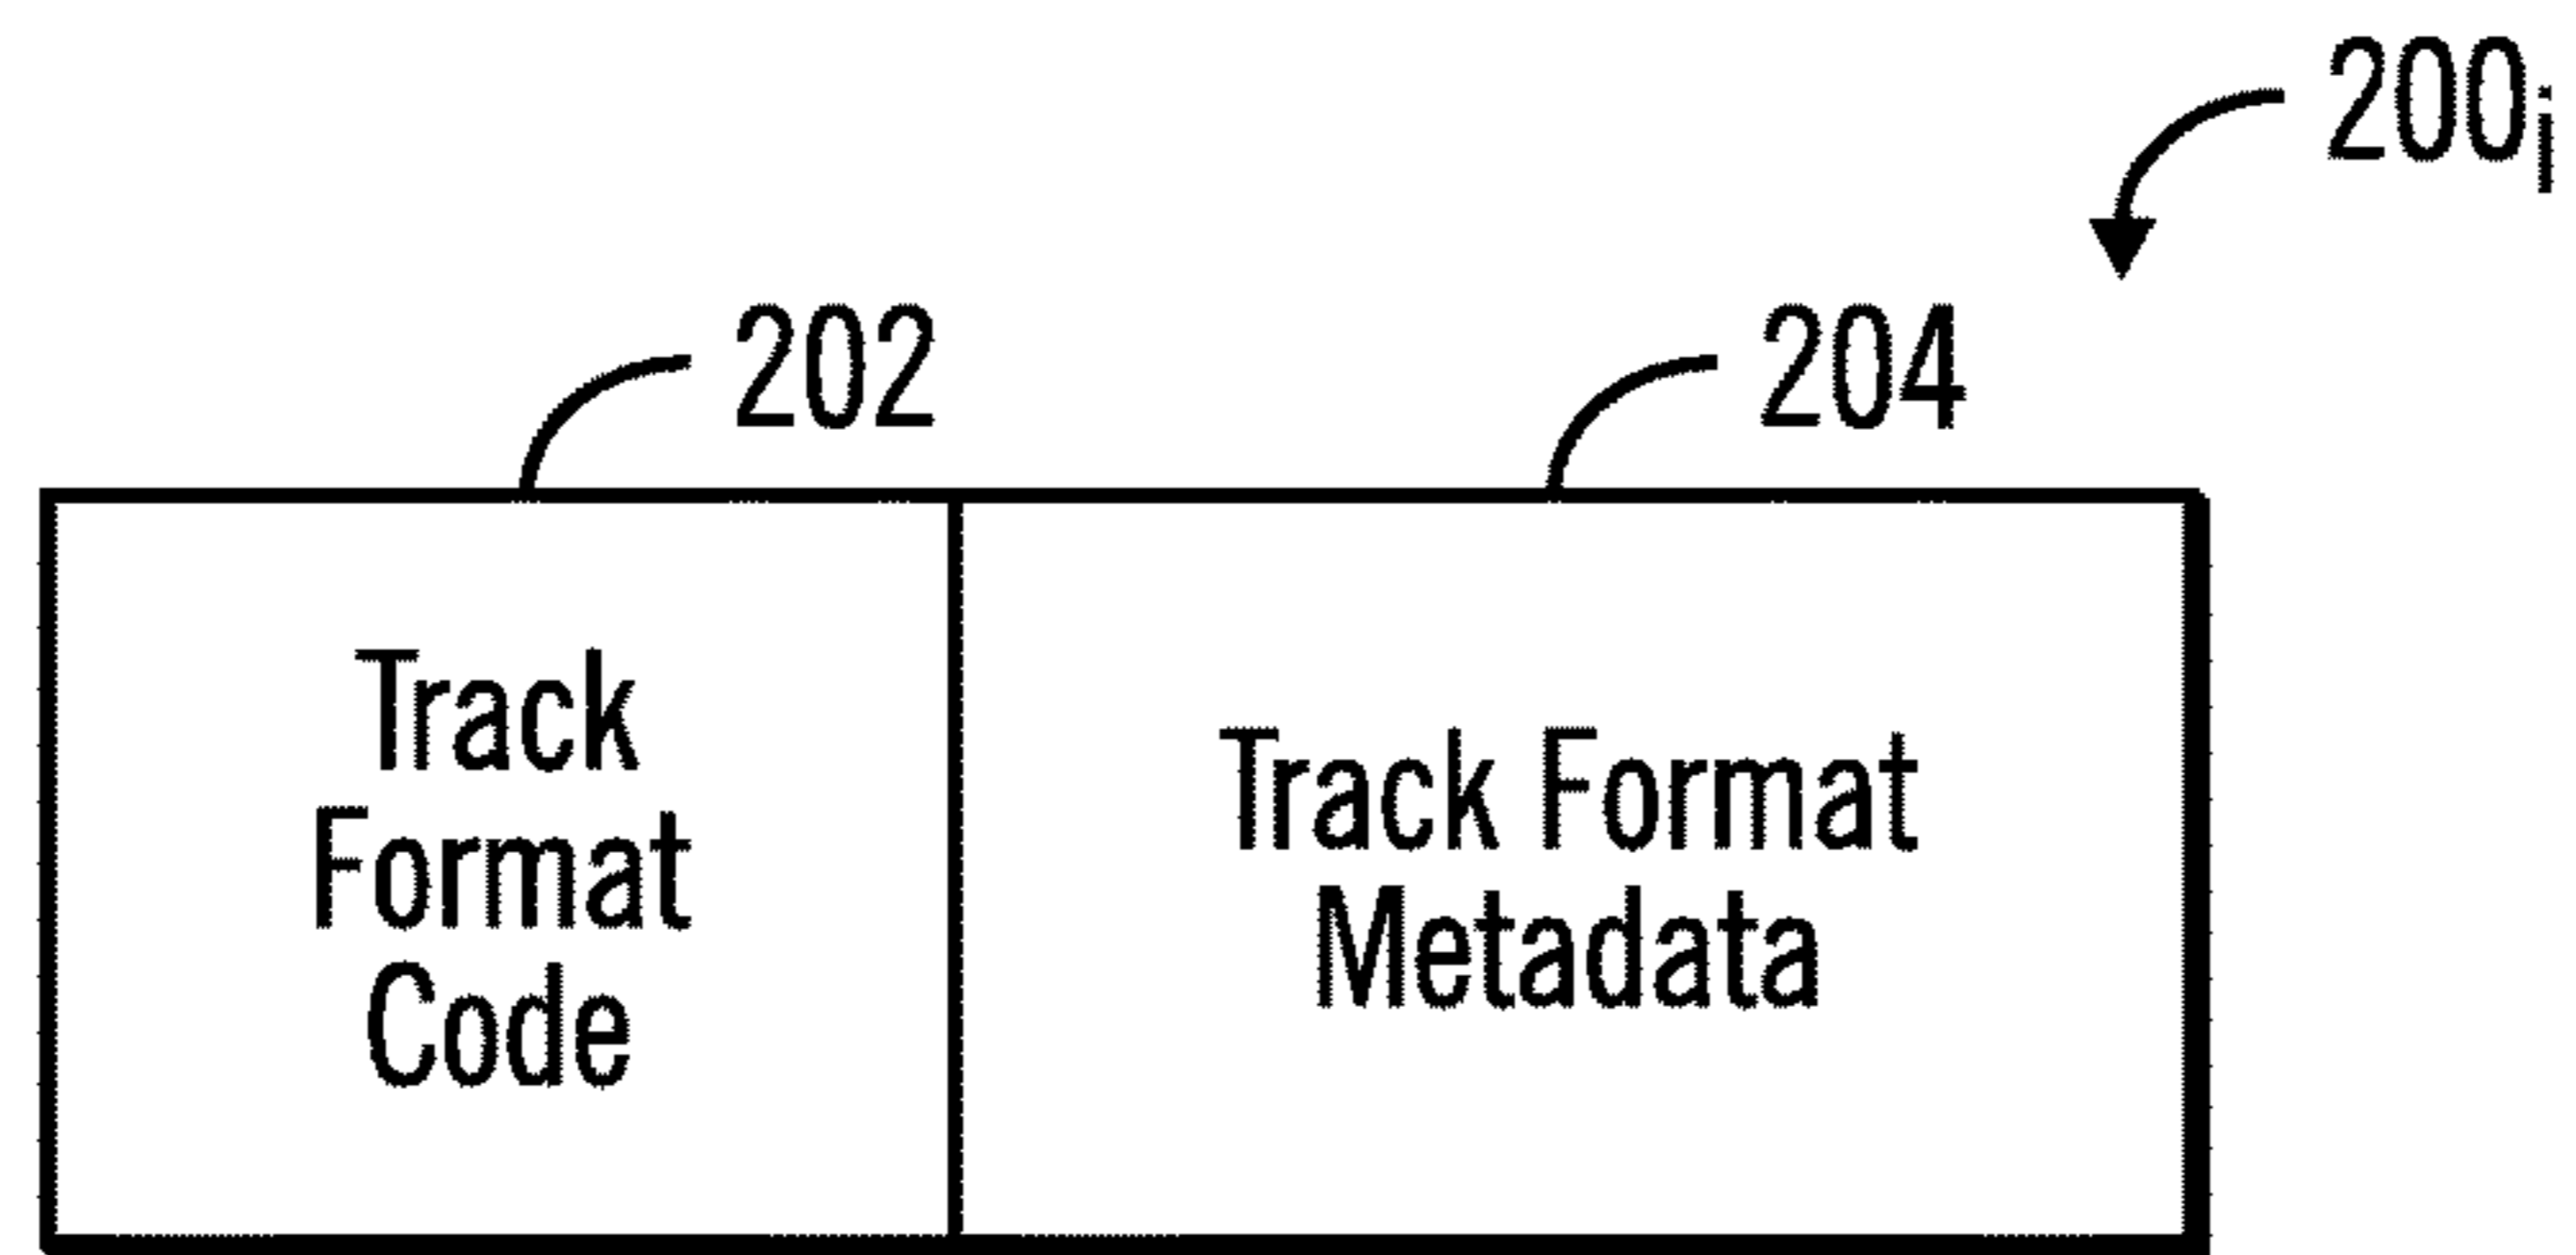
List of IBM Patents or Patent Applications Treated as Related, dated Jan. 14, 2020, 2 pages [18.754C1 (Appendix P)].
Final Office Action dated Sep. 12, 2019, pp. 19, for U.S. Appl. No. 15/662,194, (18.733).
Notice of Allowance 2 for U.S. Appl. No. 15/662,200, dated Jan. 10, 2020, 13 pp. [18.734 (NOA2)].
Office Action dated Oct. 12-18, 2019, pp. 40, for U.S. Appl. No. 15/672,470, (18.754).
Response dated Jan. 14, 2019, pp. 10, to Office Action dated Oct. 12-18, 2019, pp. 40, for U.S. Appl. No. 15/672,470, (18.754).
Supplemental Amendment dated Jan. 25, 2019, pp. 8, to Office Action dated Oct. 12-18, 2019, pp. 40, for U.S. Appl. No. 15/672,470, (18.754).
Final Office Action dated May 9, 2019, pp. 38, for U.S. Appl. No. 15/672,470, (18.754).
Response dated Jul. 11, 2019, pp. 12, to Final Office Action dated May 9, 2019, pp. 38, for U.S. Appl. No. 15/672,470, (18.754).
Notice of Allowance dated Sep. 5, 2019, pp. 22, for U.S. Appl. No. 15/672,470, (18.754).
Notice of Allowance 2 for U.S. Appl. No. 15/672,470, dated Oct. 17, 2019, 15 pp. [18.754 (NOA2)].

U.S. Appl. No. 16/836,827, filed Mar. 31, 2020, (18.733C1).
Preliminary Amendment dated Mar. 31, 2020, pp. 9,(18.733C1).
Amendment filed Mar. 10, 2020, pp. 12, to Final Office Action dated Dec. 31, 2019, pp. 57, for U.S. Appl. No. 15/672,460, (18.735).
Notice of Allowance dated Apr. 8, 2020, pp. 14, for U.S. Appl. No. 15/672,460, (18.735).
List of IBM Patents and Applications Treated as Related, dated Apr. 9, 2020, pp. 2.
Amendment dated Mar. 16, 2020, pp. 14, for U.S. Appl. No. 15/662,200, (18.734).
Office Action dated 2020-20-20, pp. 18, for U.S. Appl. No. 15/662,200, (18.734).
Office Action dated Oct. 2, 2020, pp. 15, for U.S. Appl. No. 15/672,460, (18.735).
Response dated Jan. 3, 2021, pp. 11, to Office Action dated Oct. 2, 2020, pp. 15, for U.S. Appl. No. 15/672,460, (18.735).
Response dated Oct. 20, 2020, pp. 15, to Office Action dated Jul. 20, 2020, pp. 18, for U.S. Appl. No. 15/662,200, (18.734).
Final Office Action dated Jan. 28, 2021, pp. 18, for U.S. Appl. No. 15/662,200, (18.734).

* cited by examiner

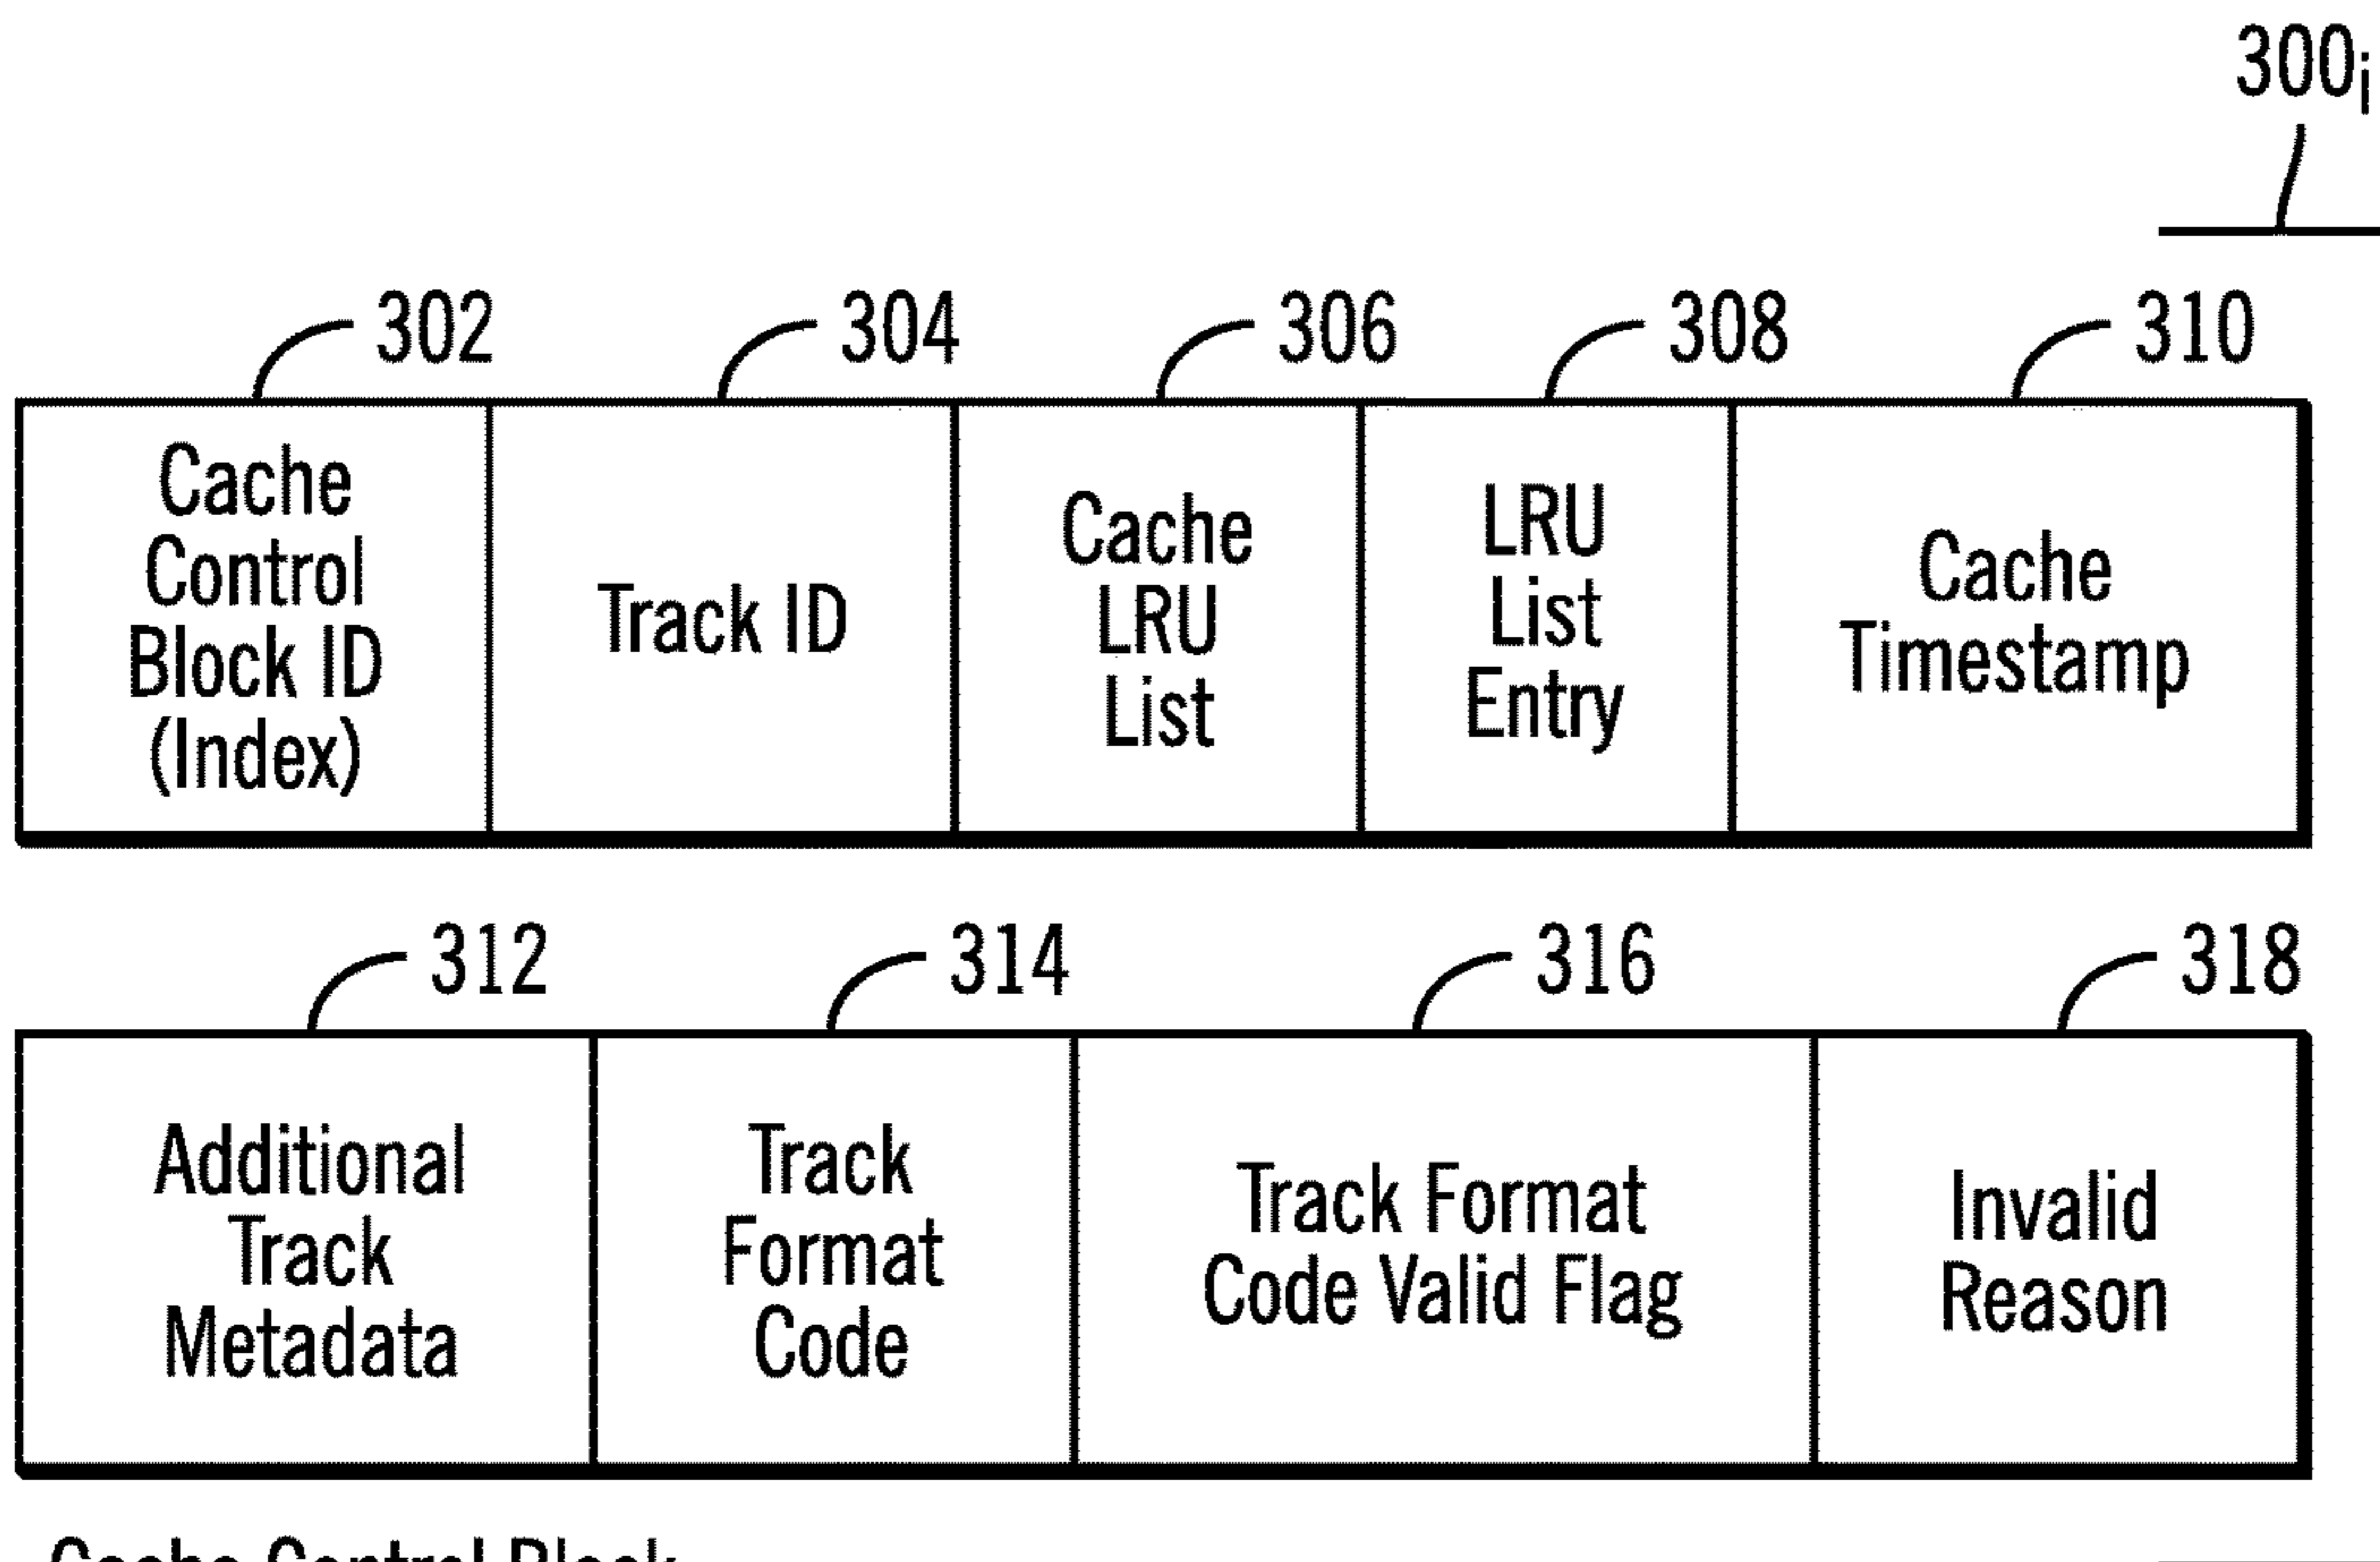
‡ imported from a related application





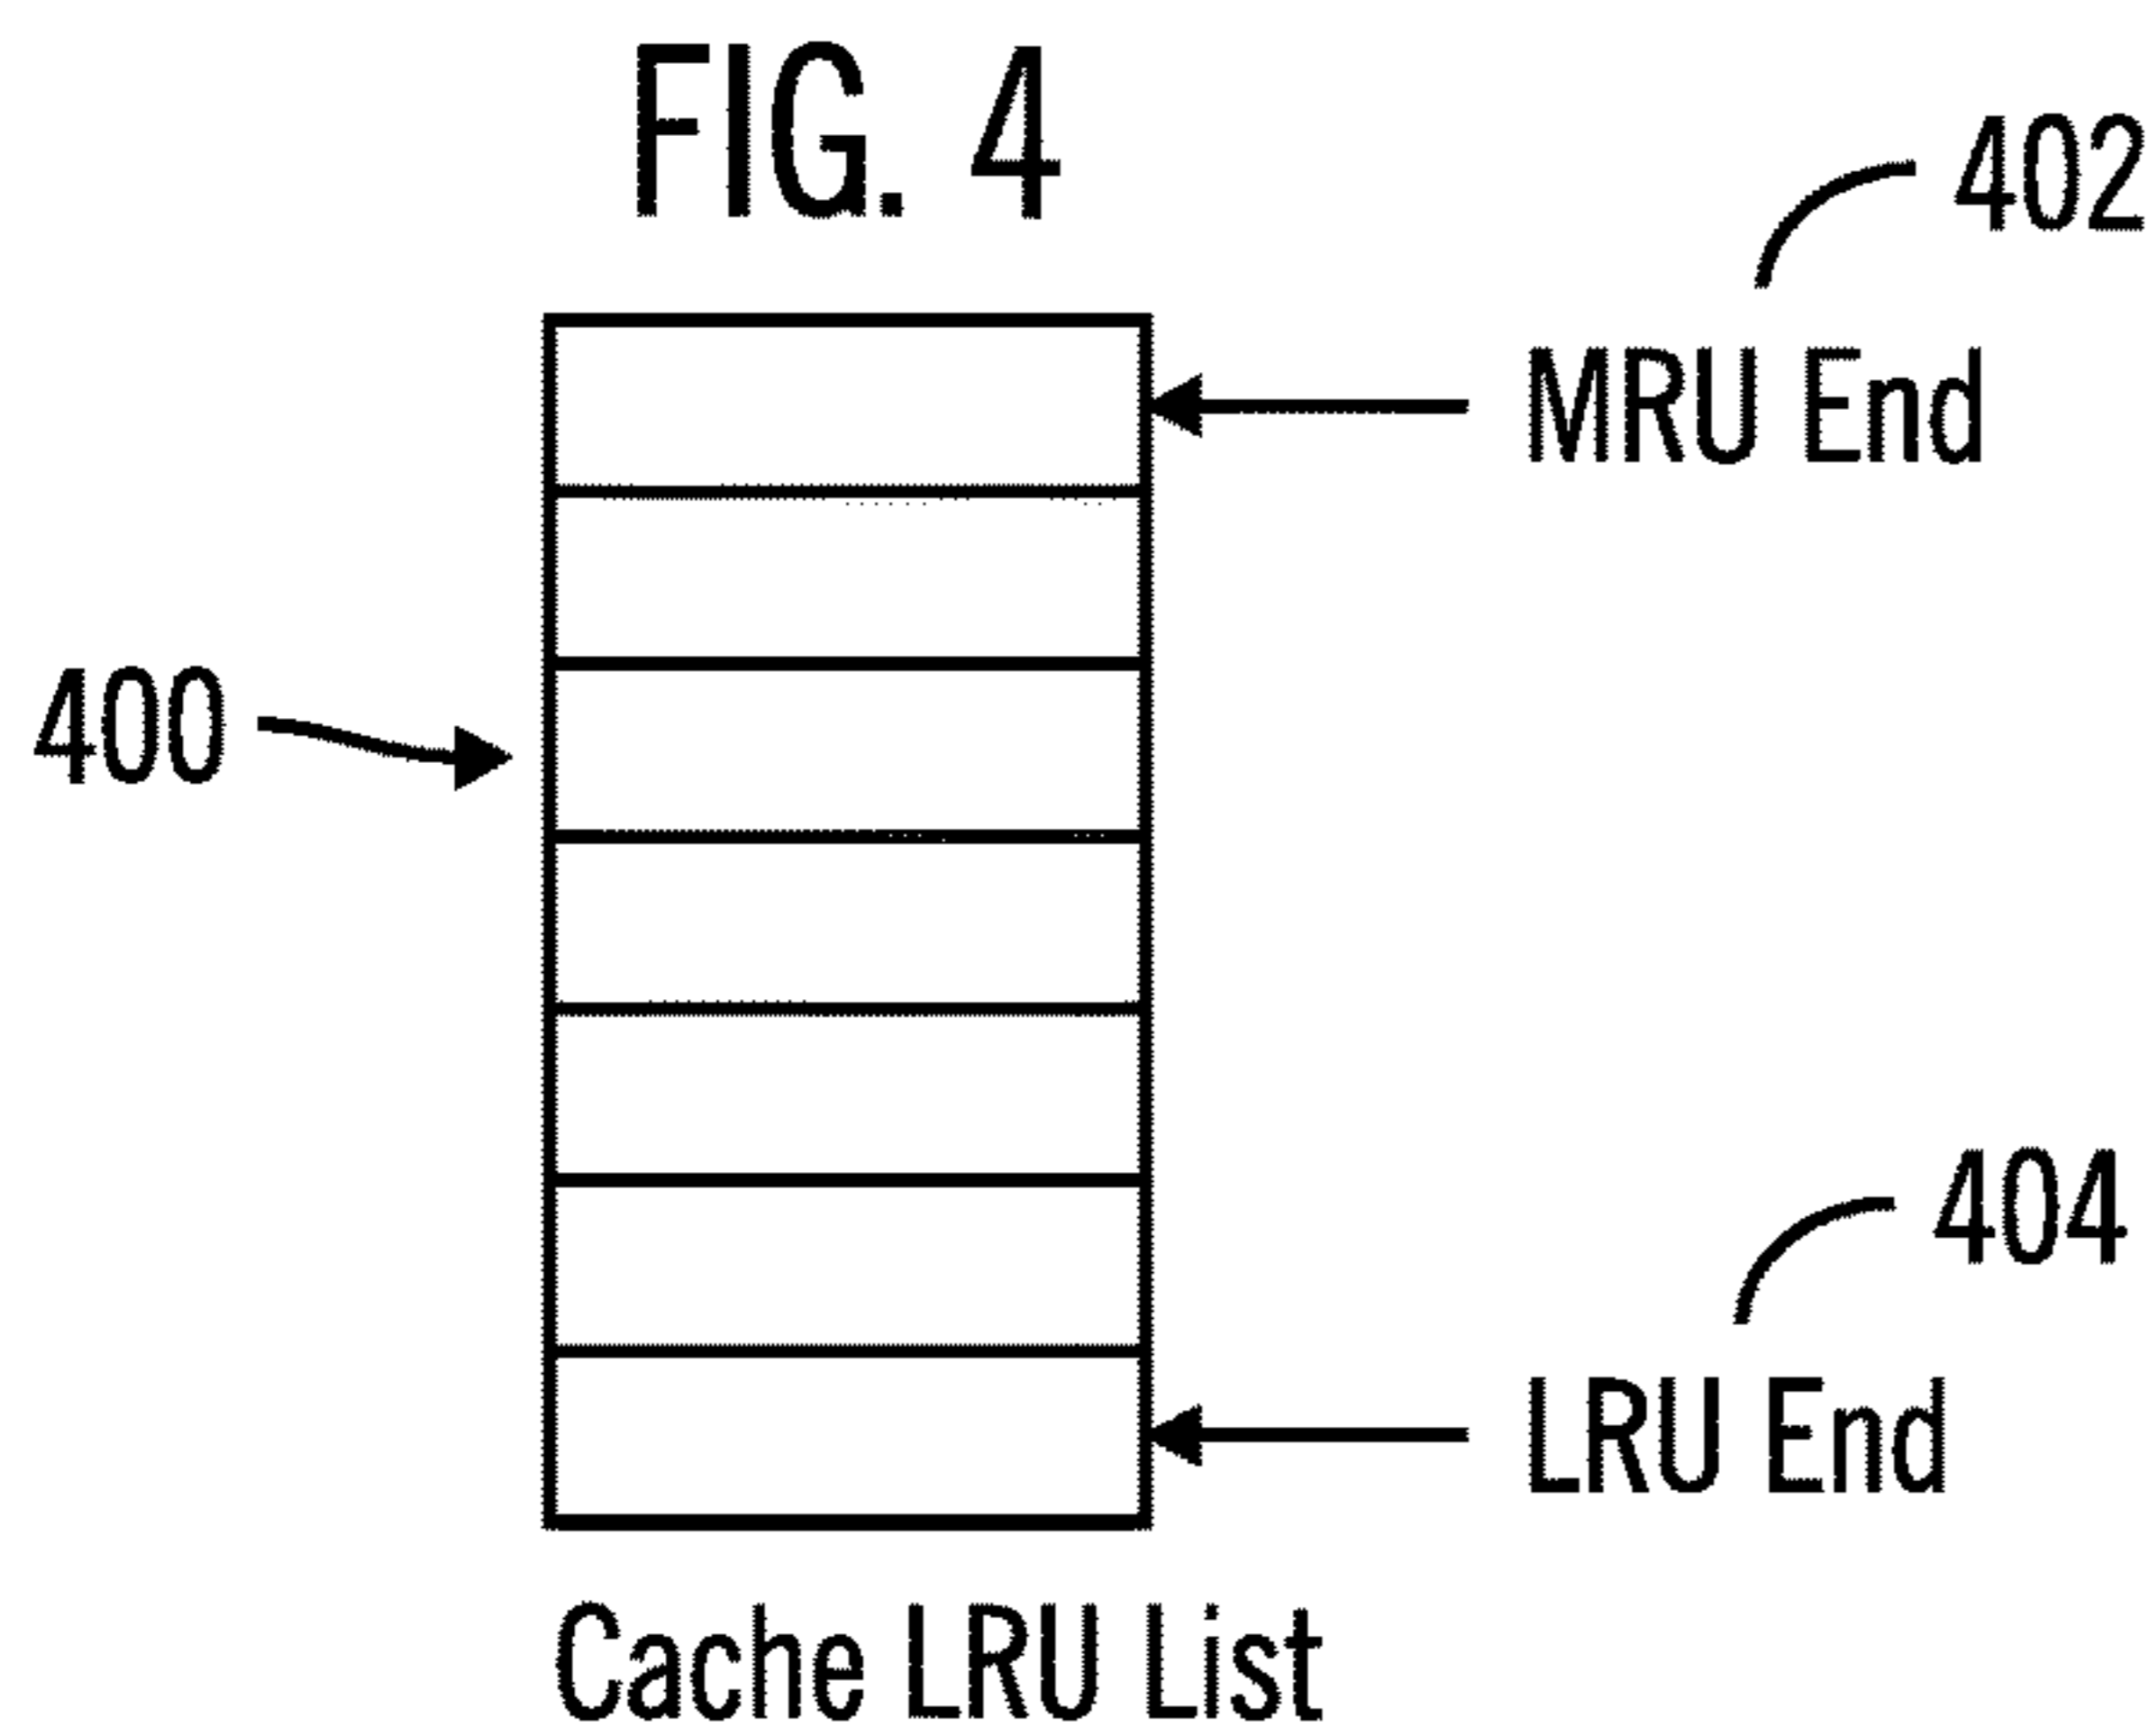
Track Format Table Entry

FIG. 2

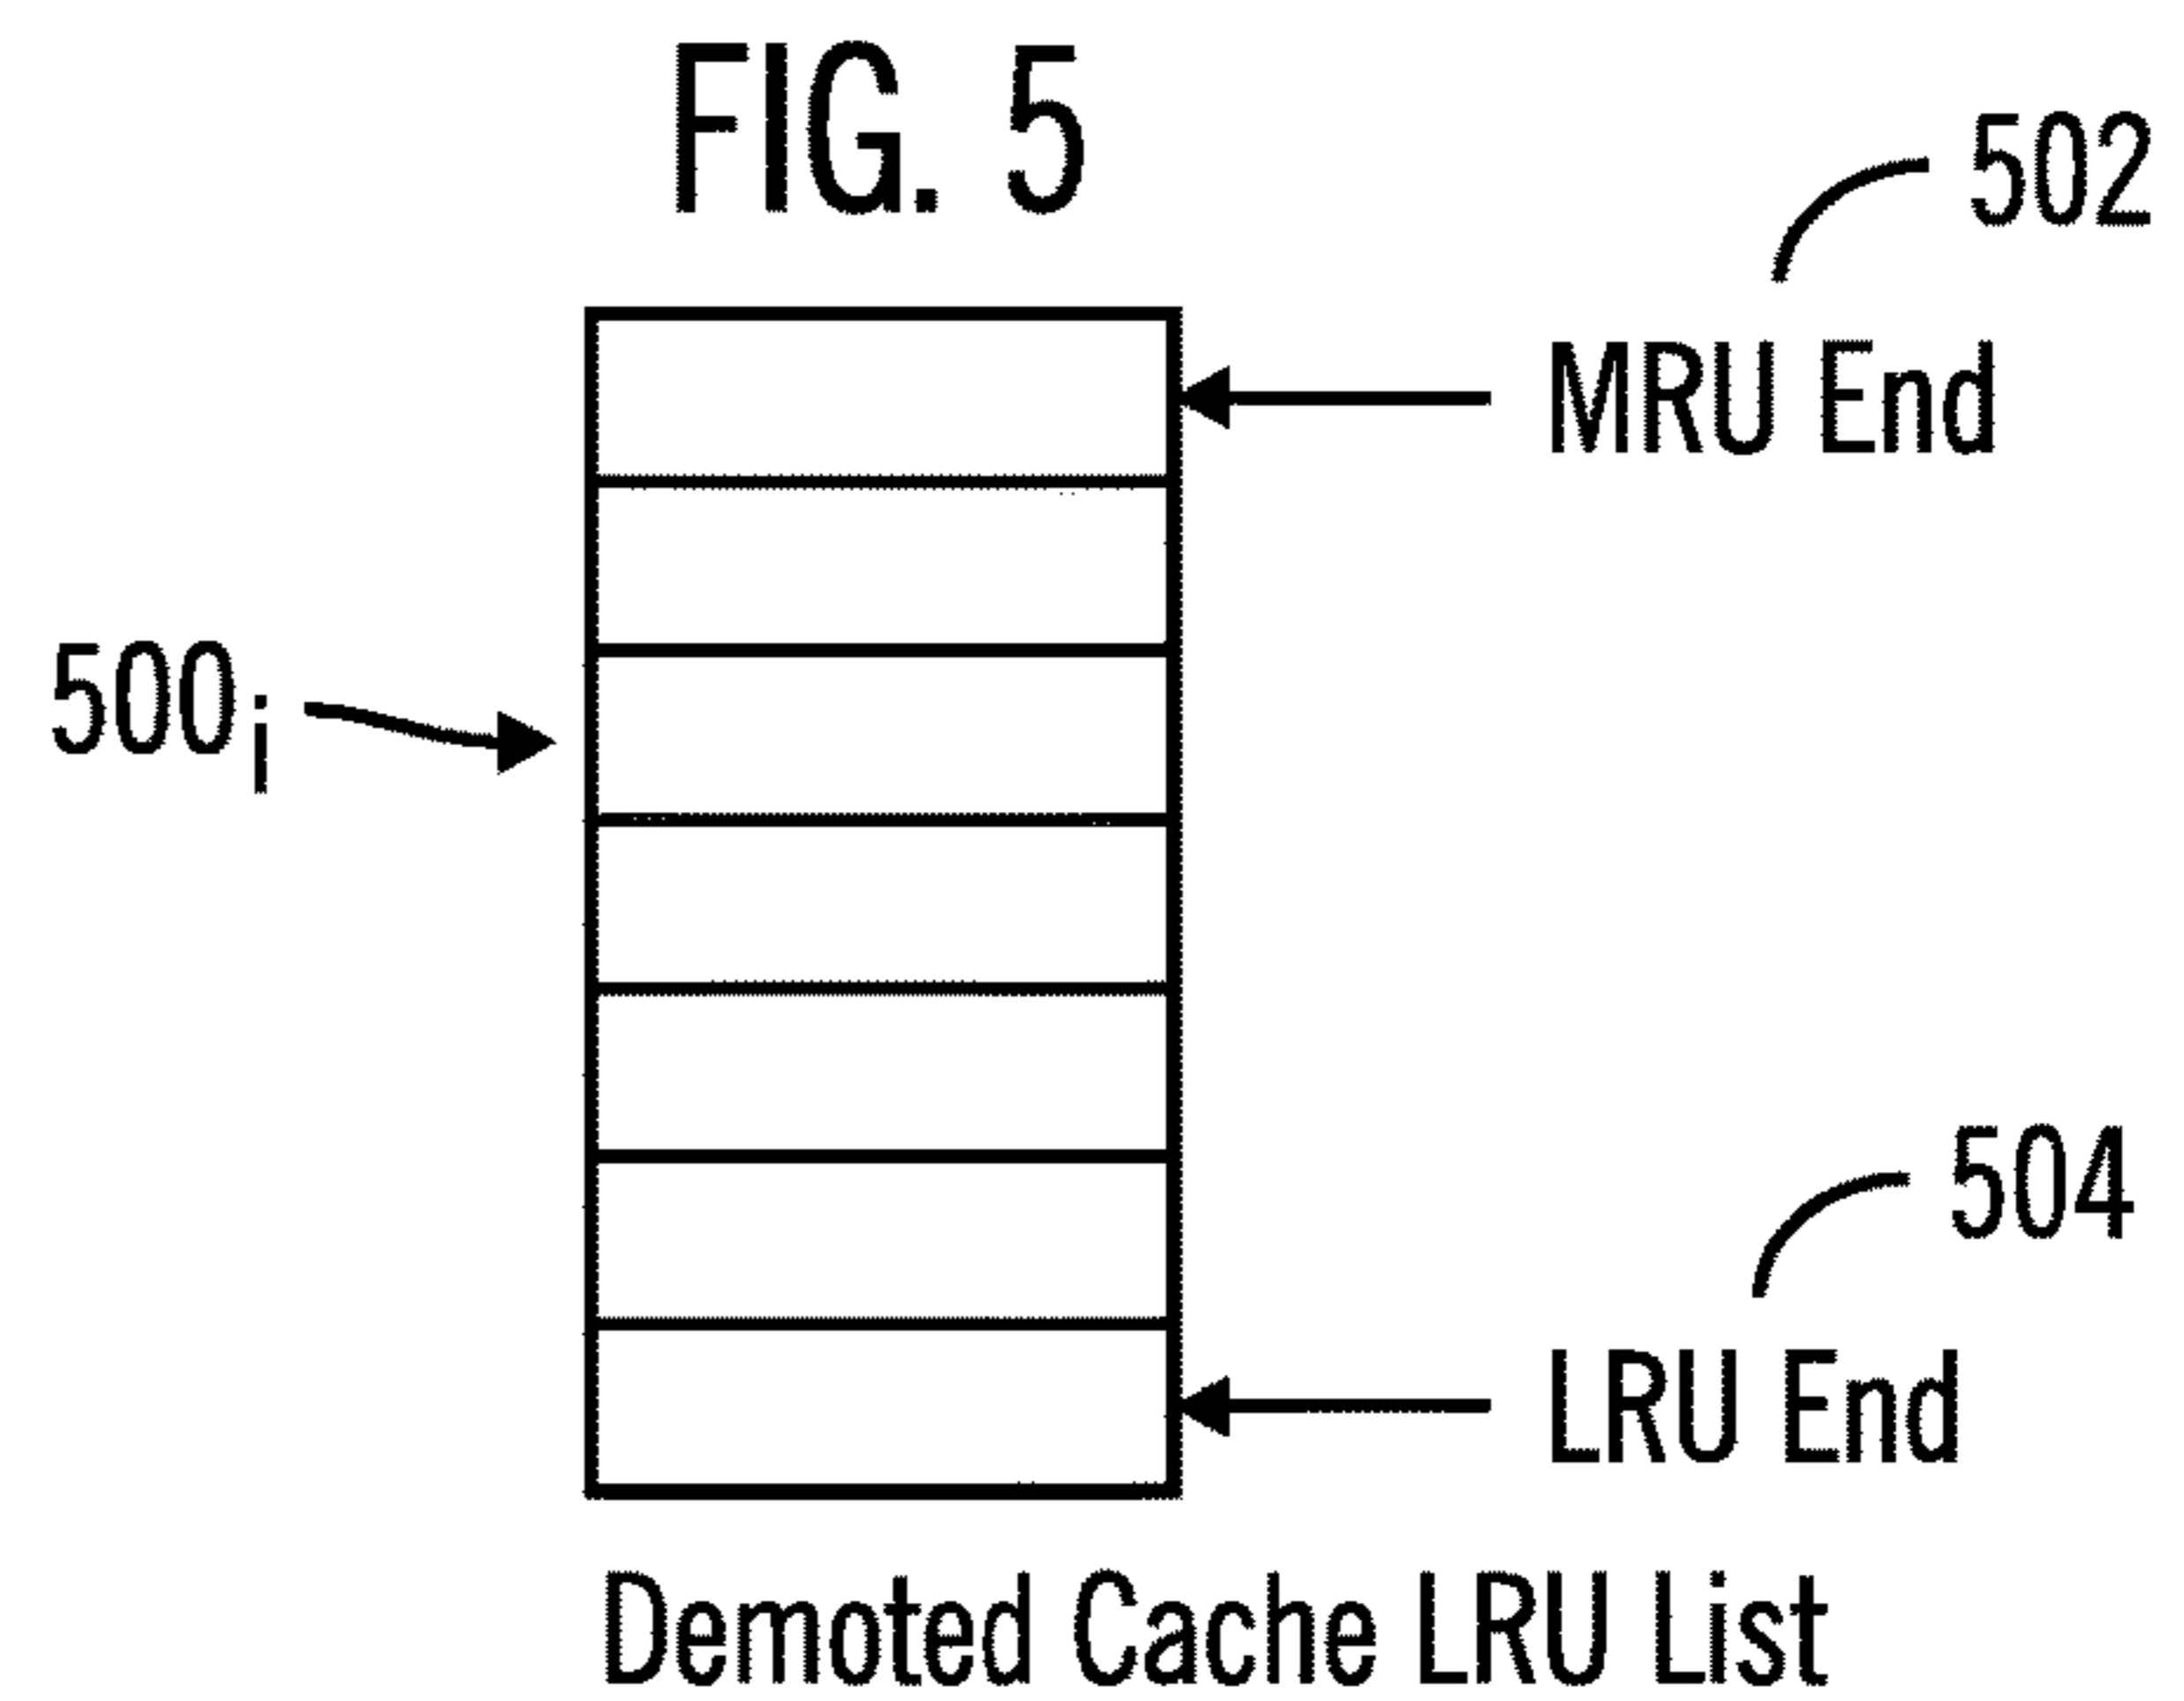


Cache Control Block

FIG. 3



Cache LRU List



Demoted Cache LRU List

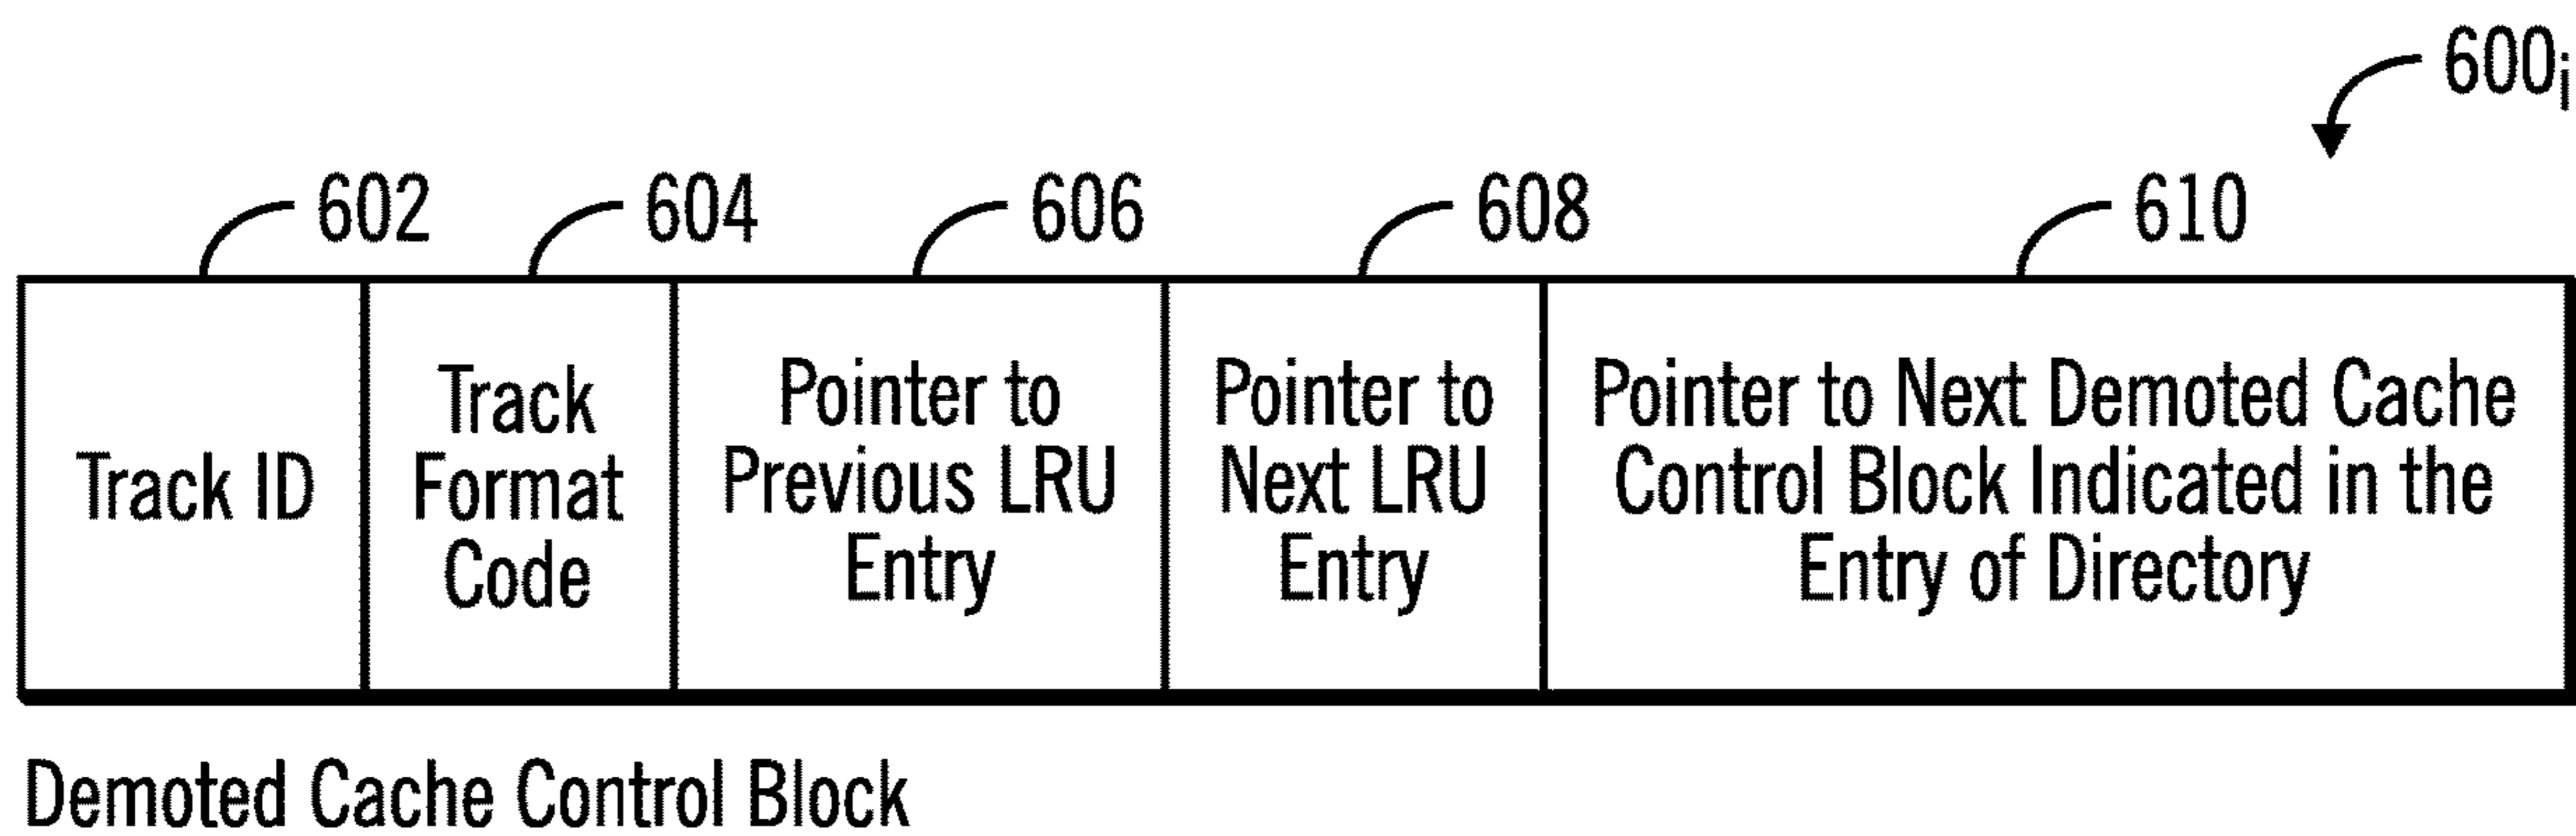


FIG. 6

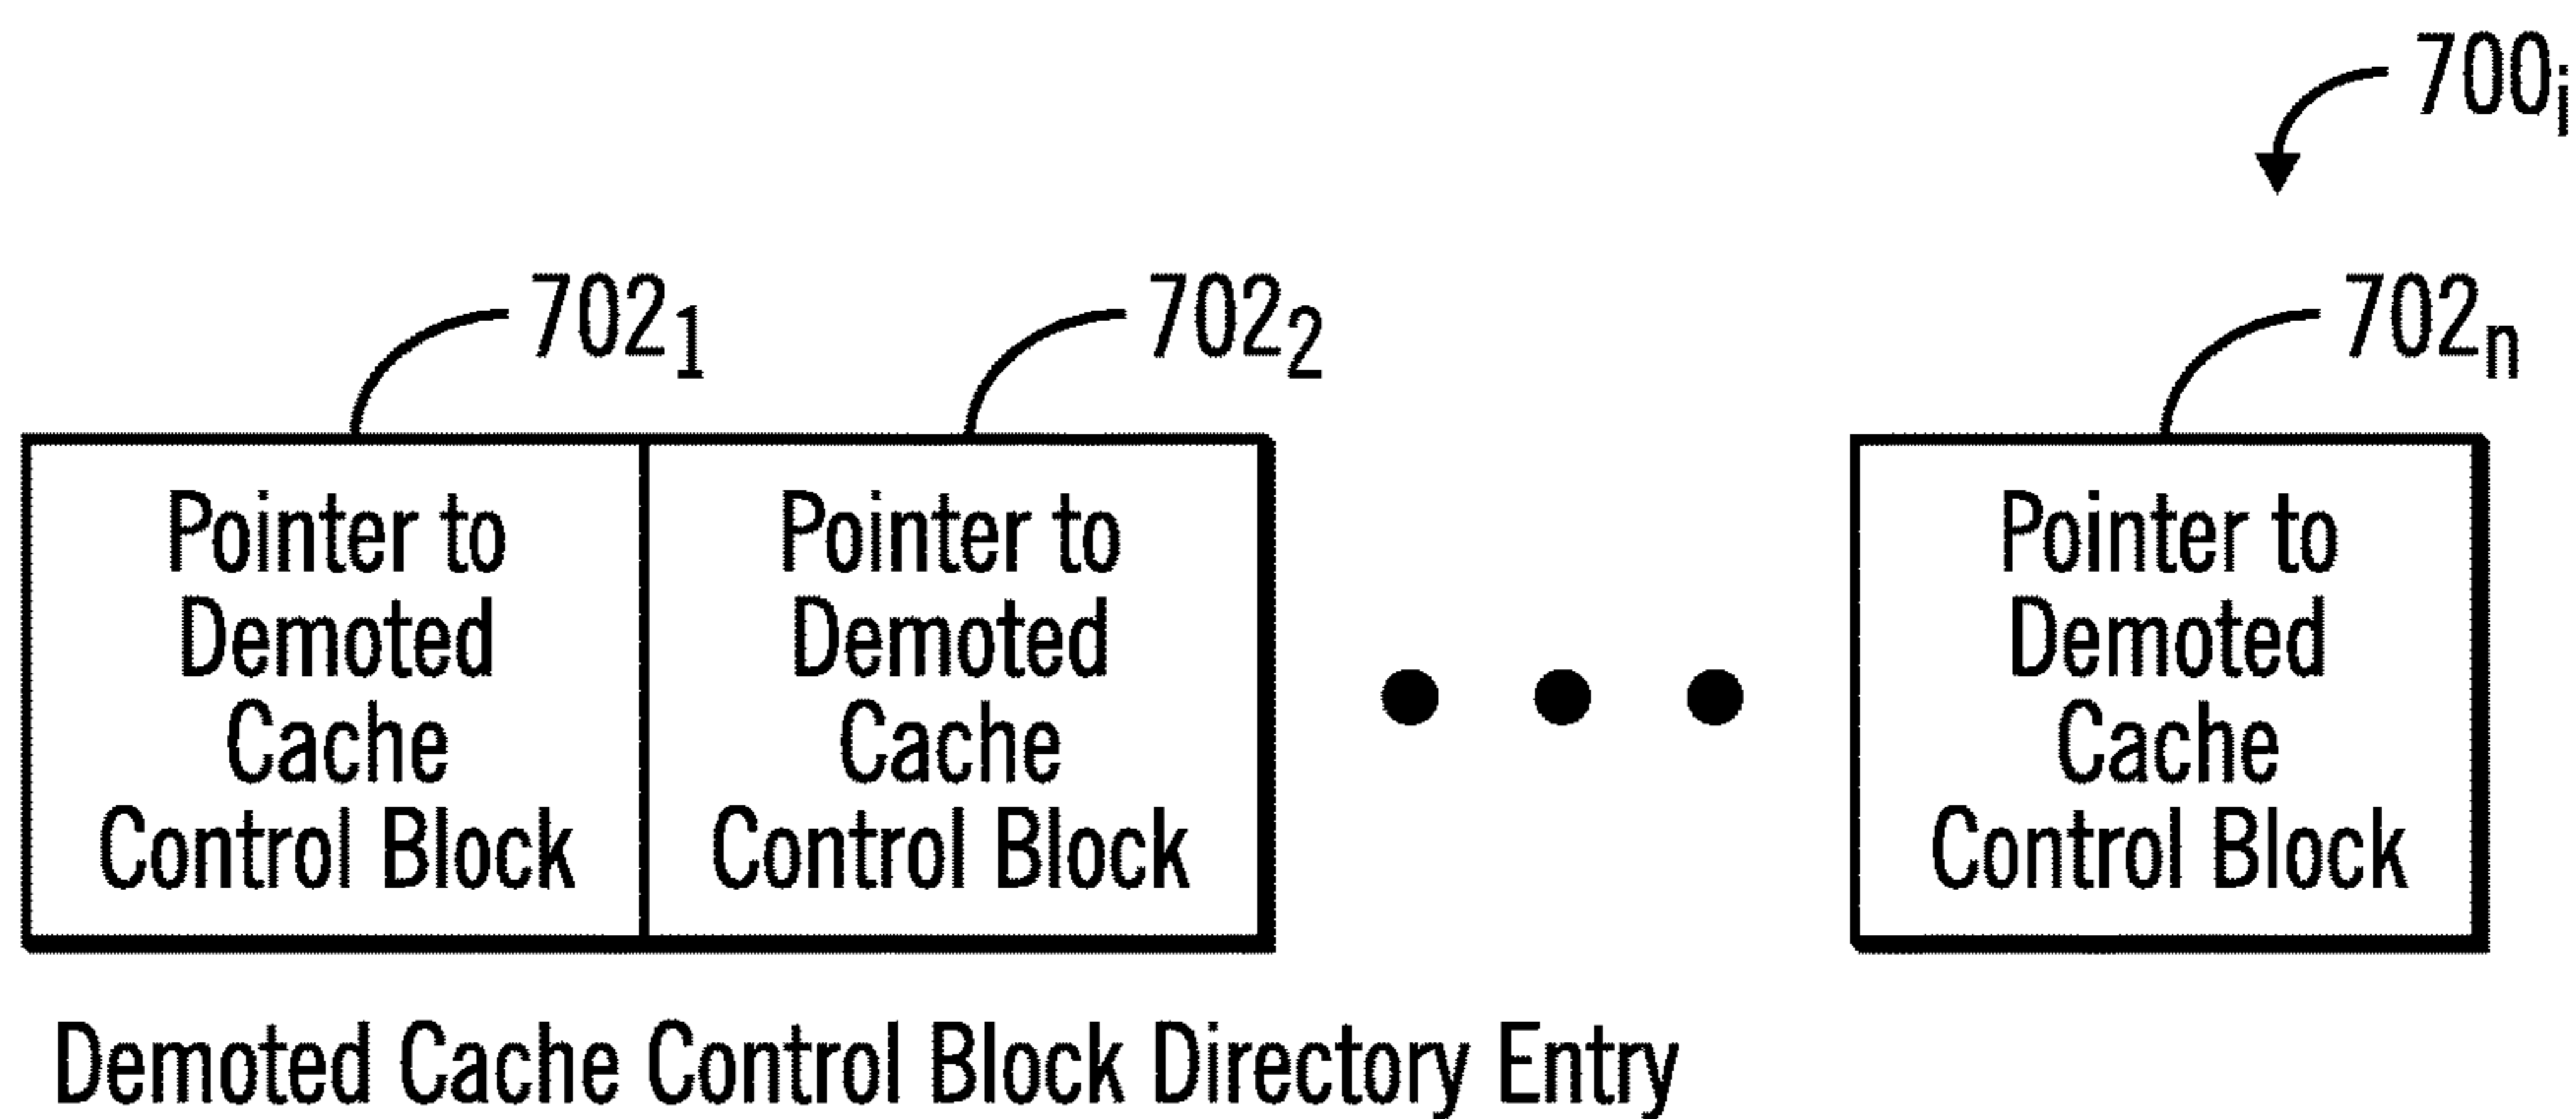


FIG. 7

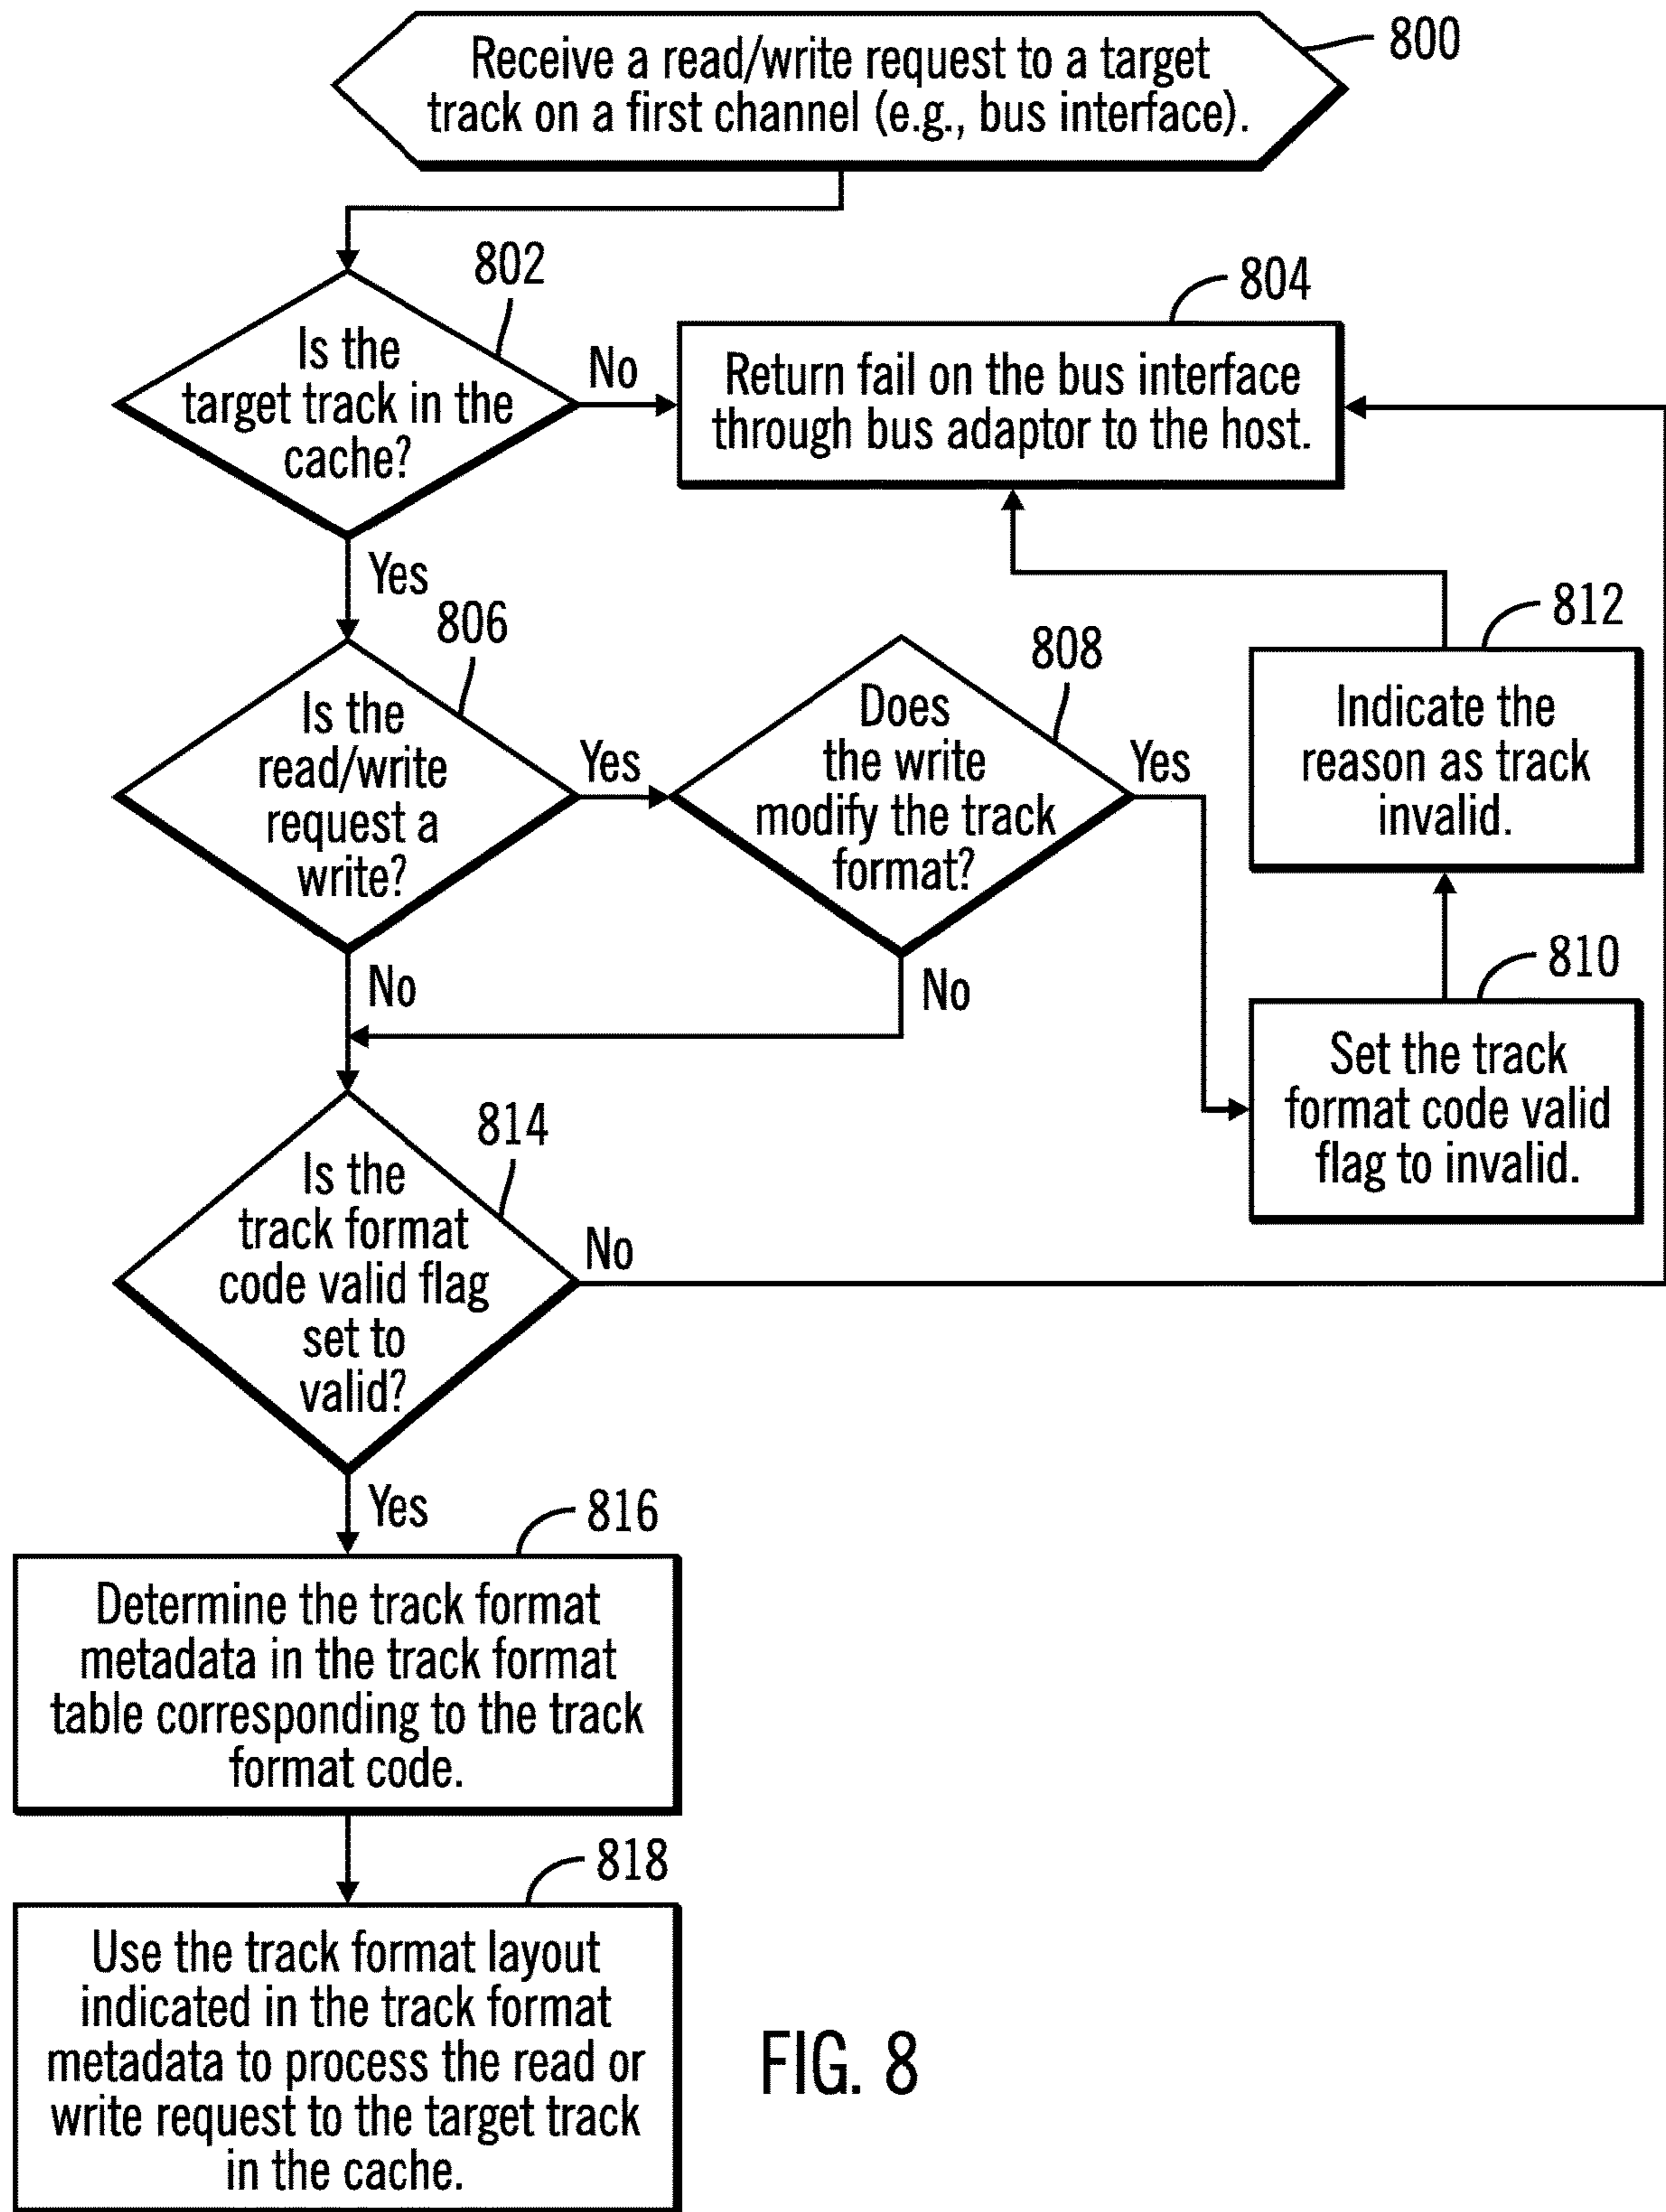


FIG. 8

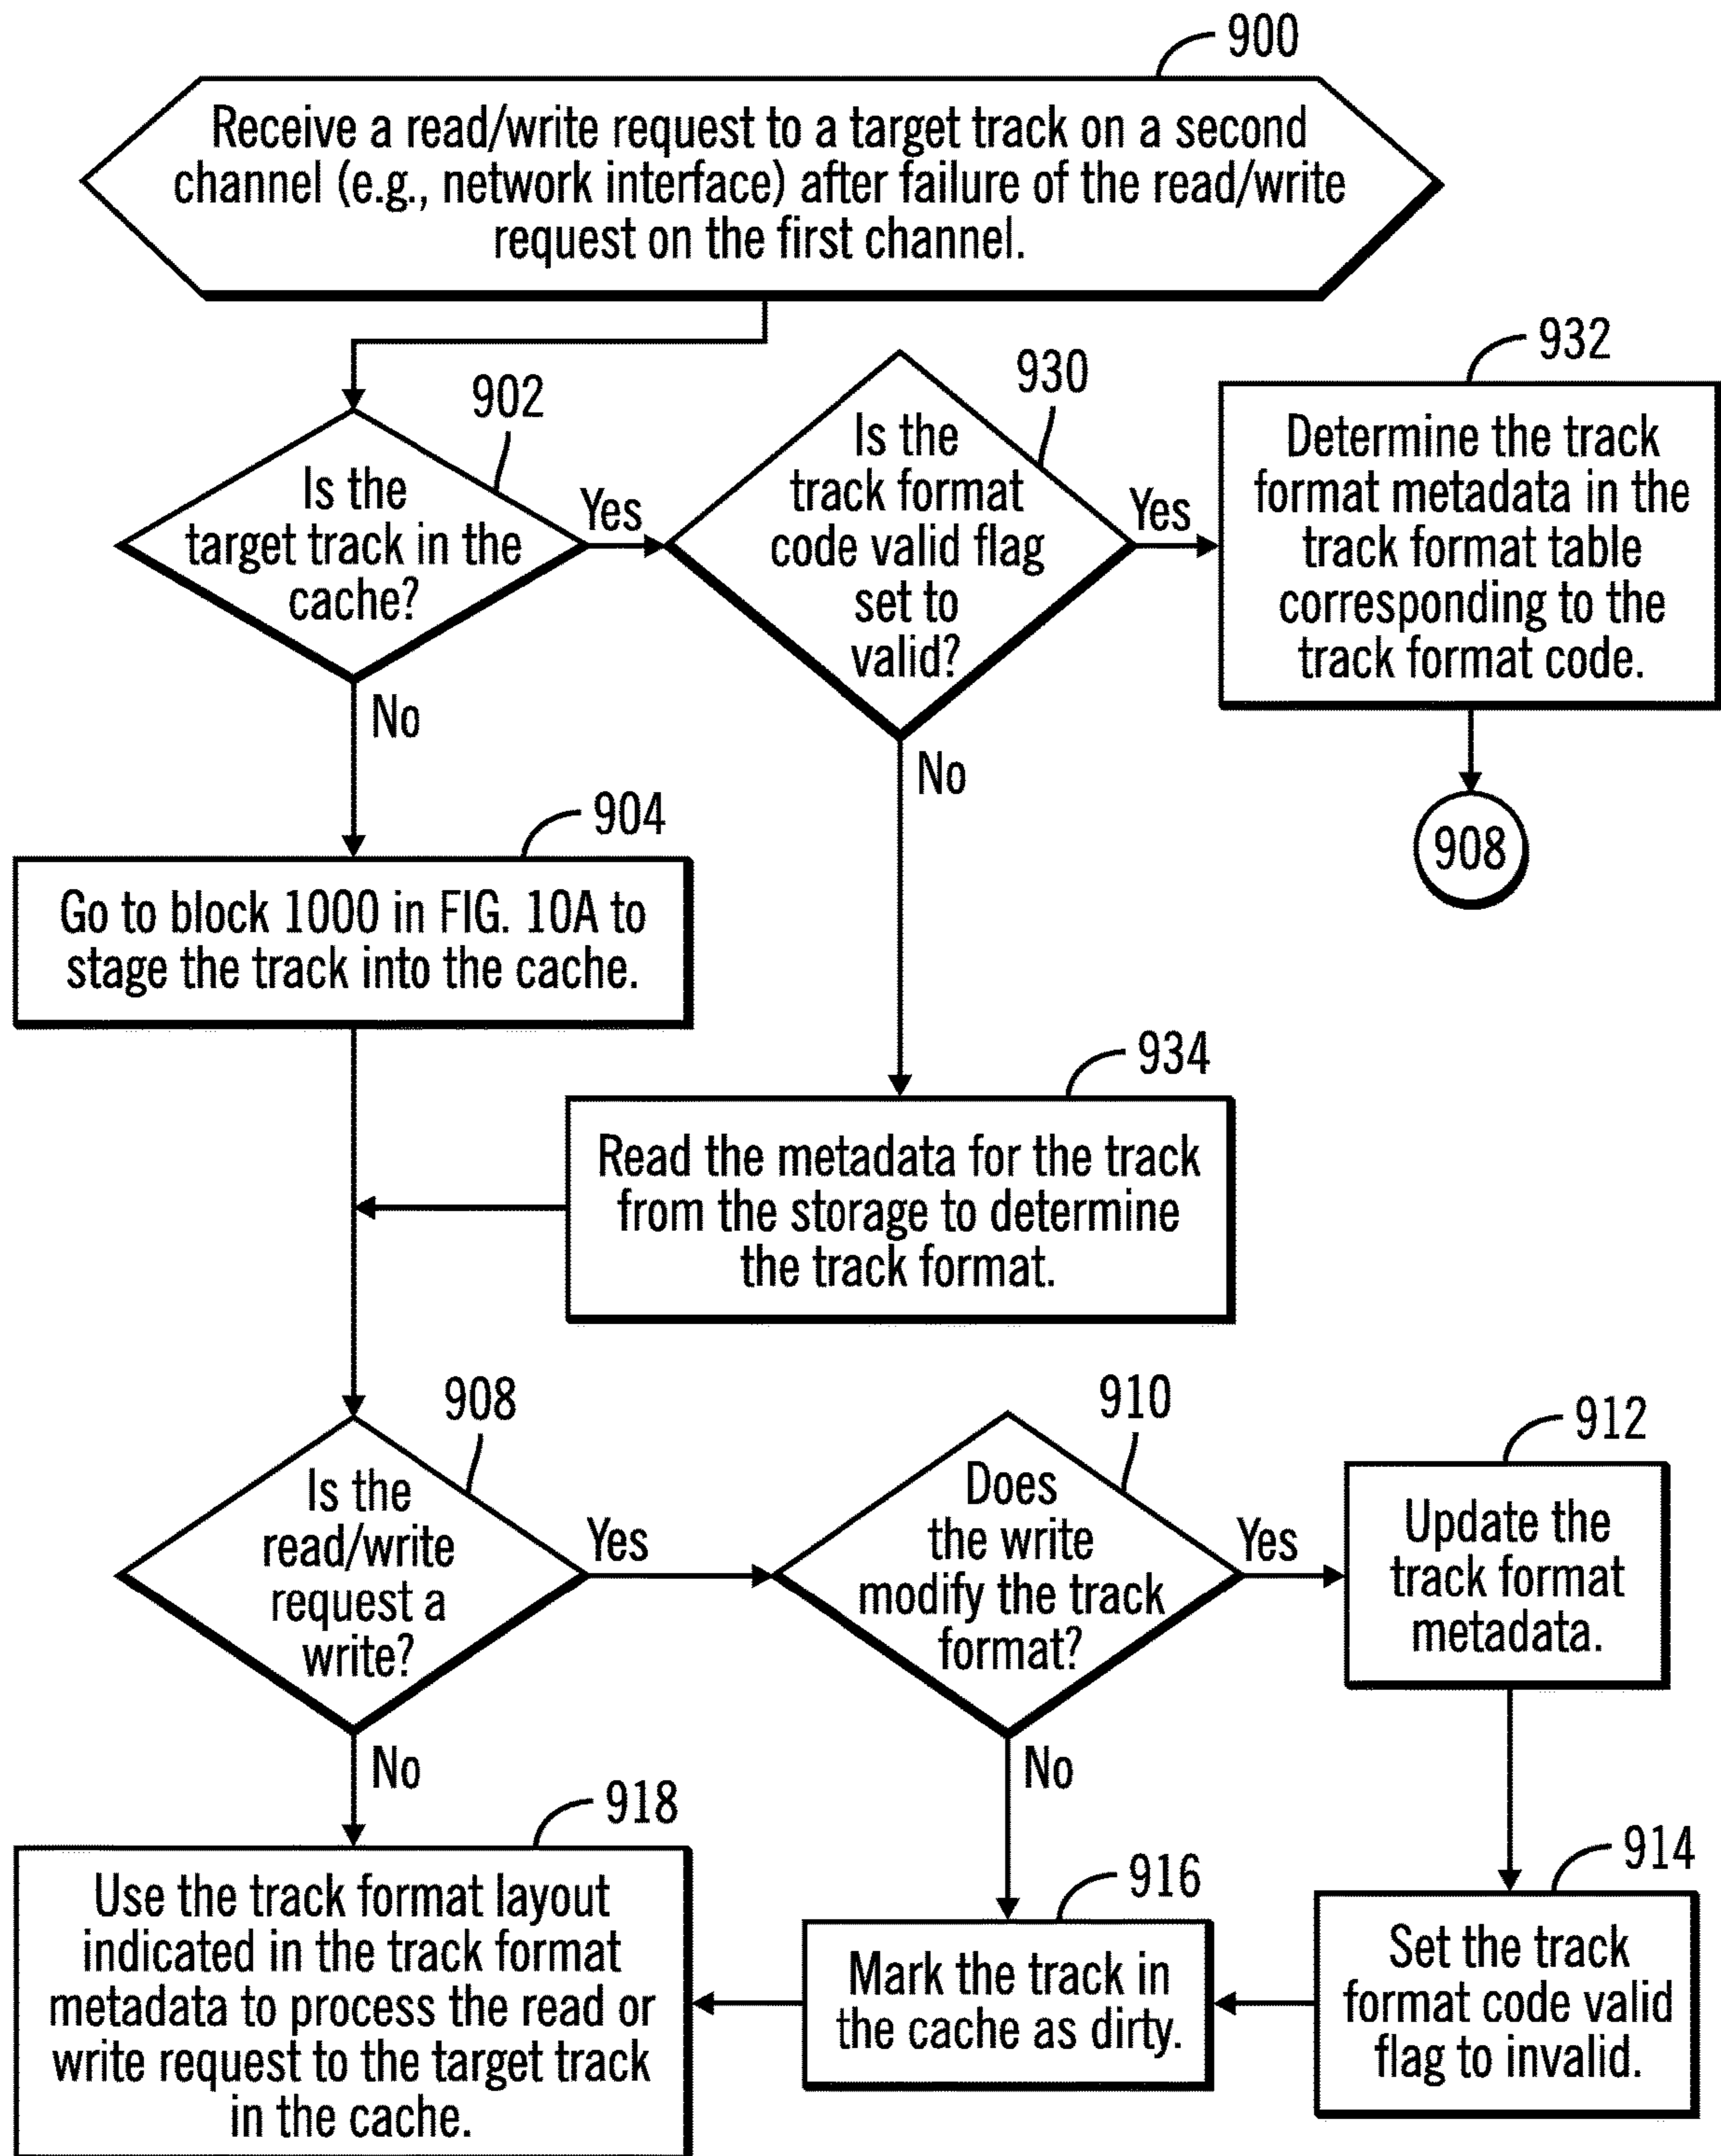
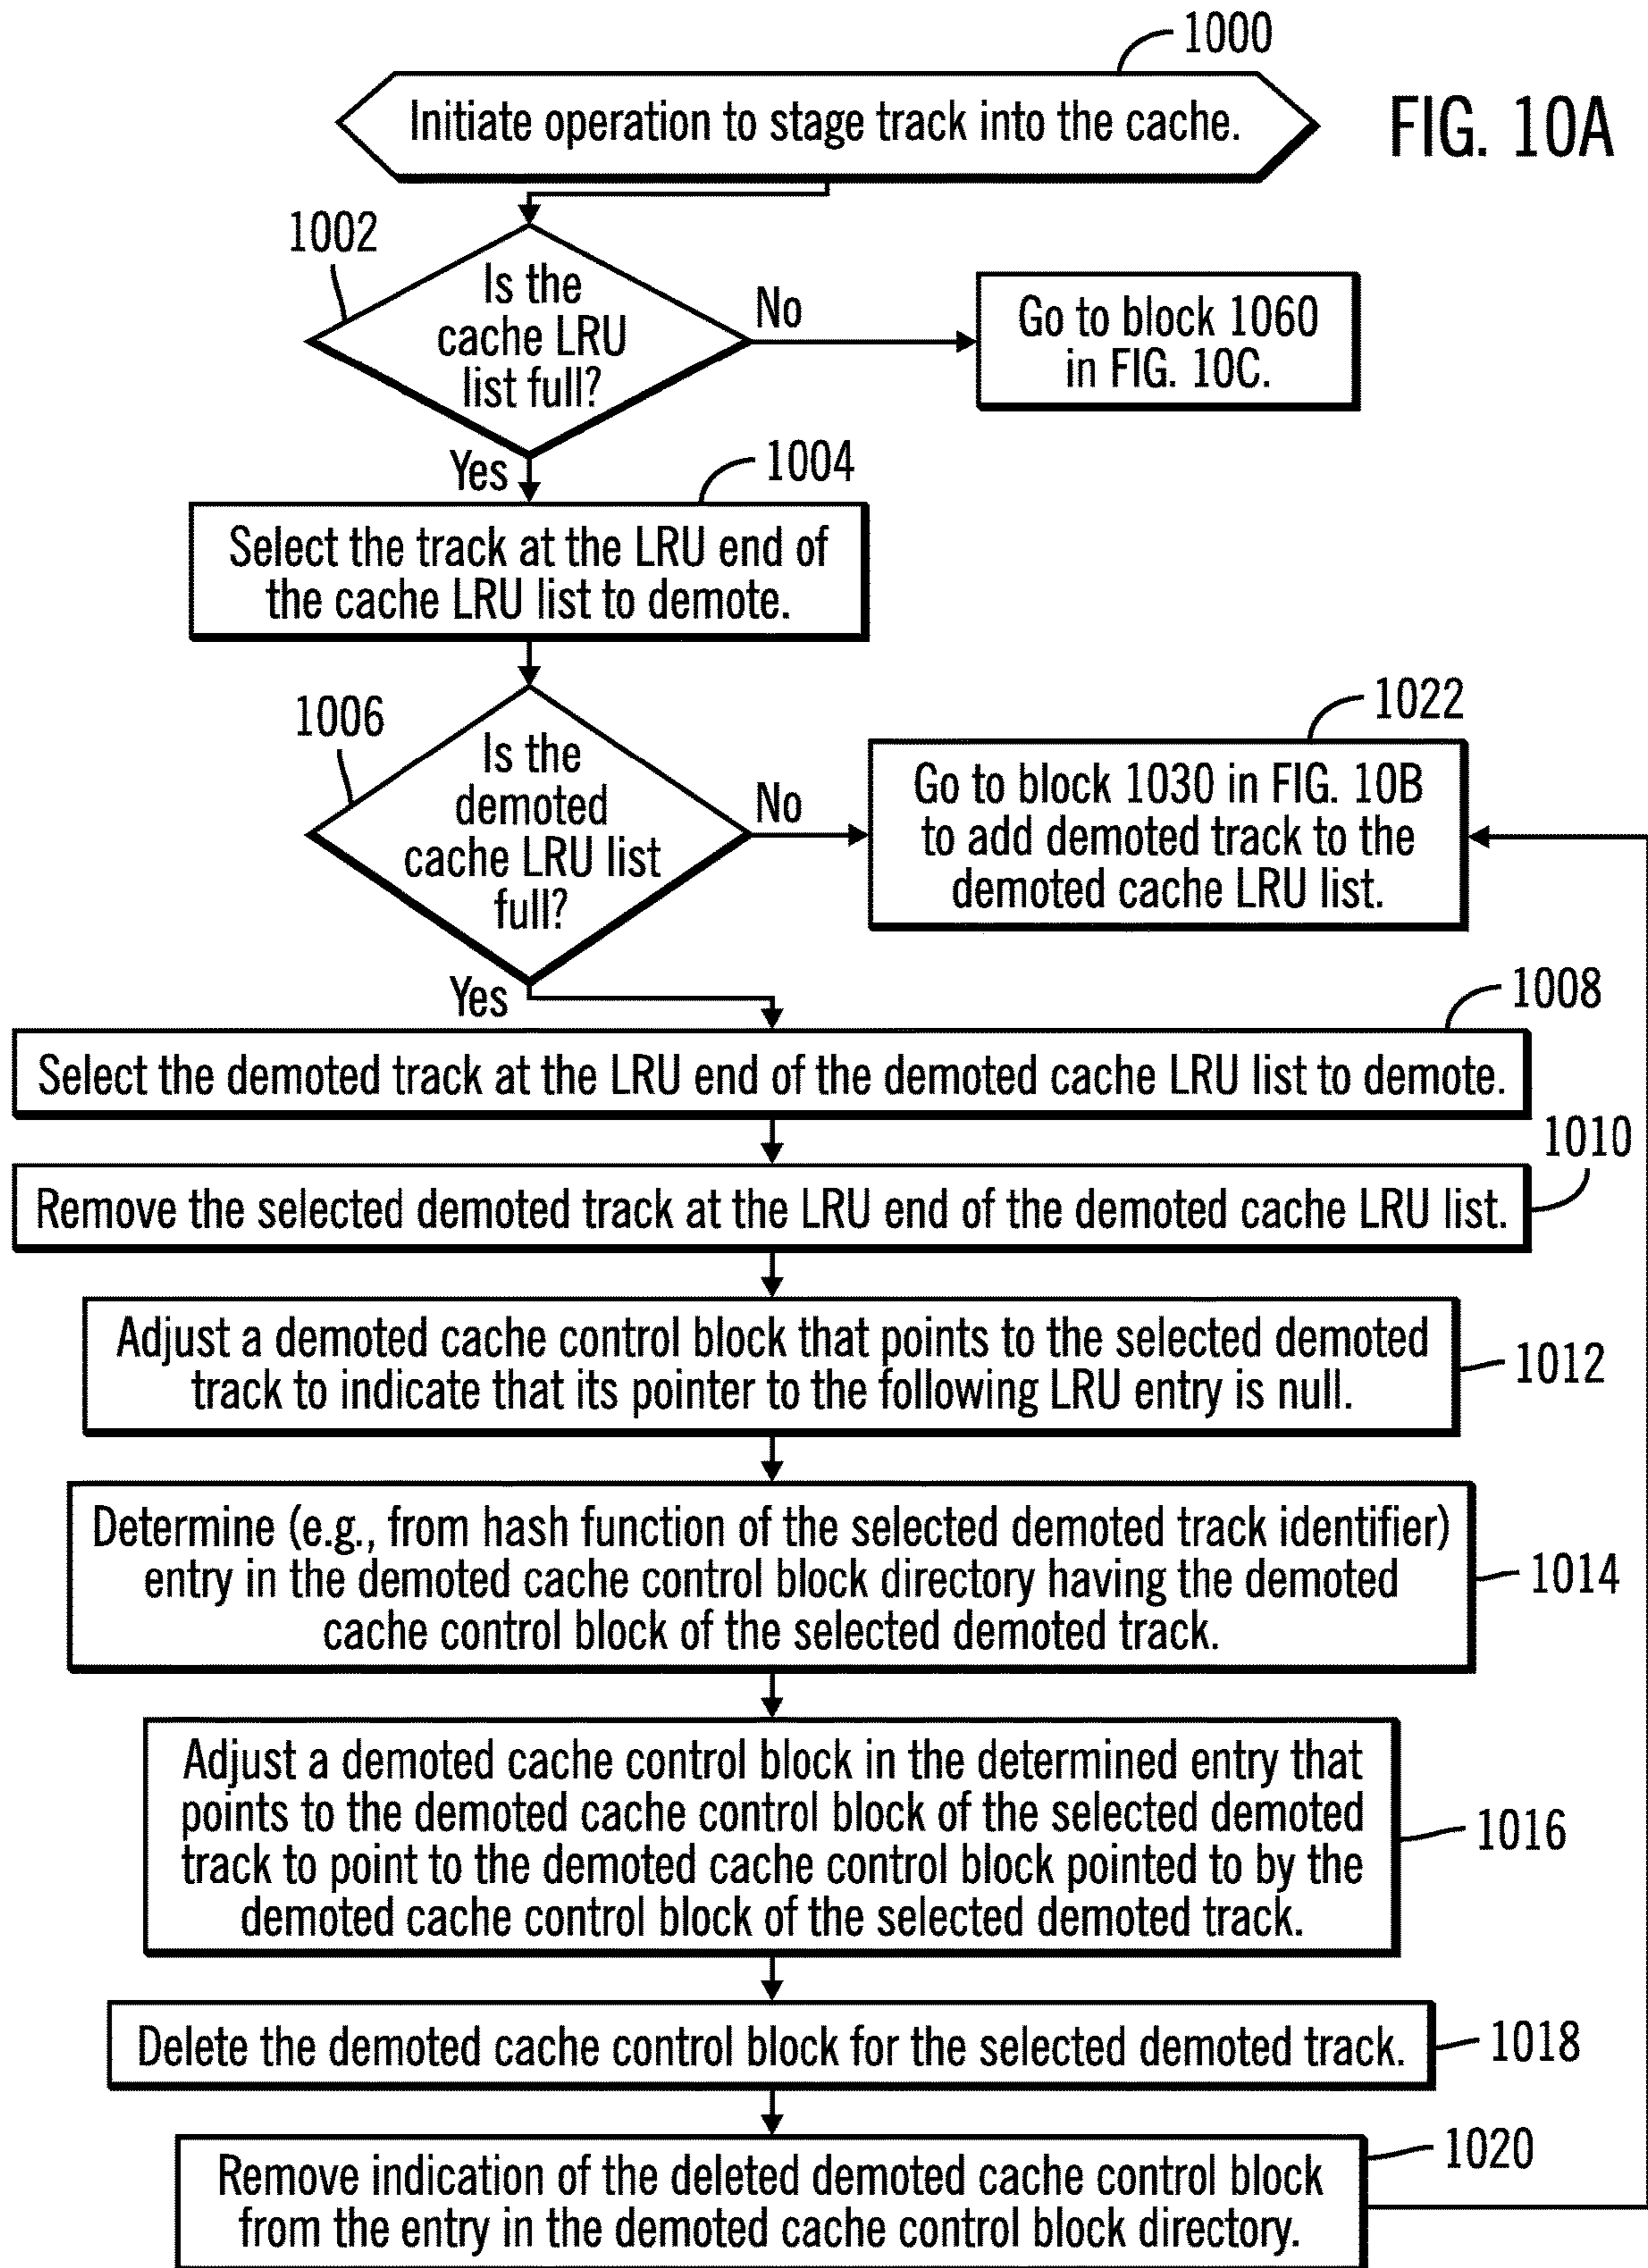


FIG. 9



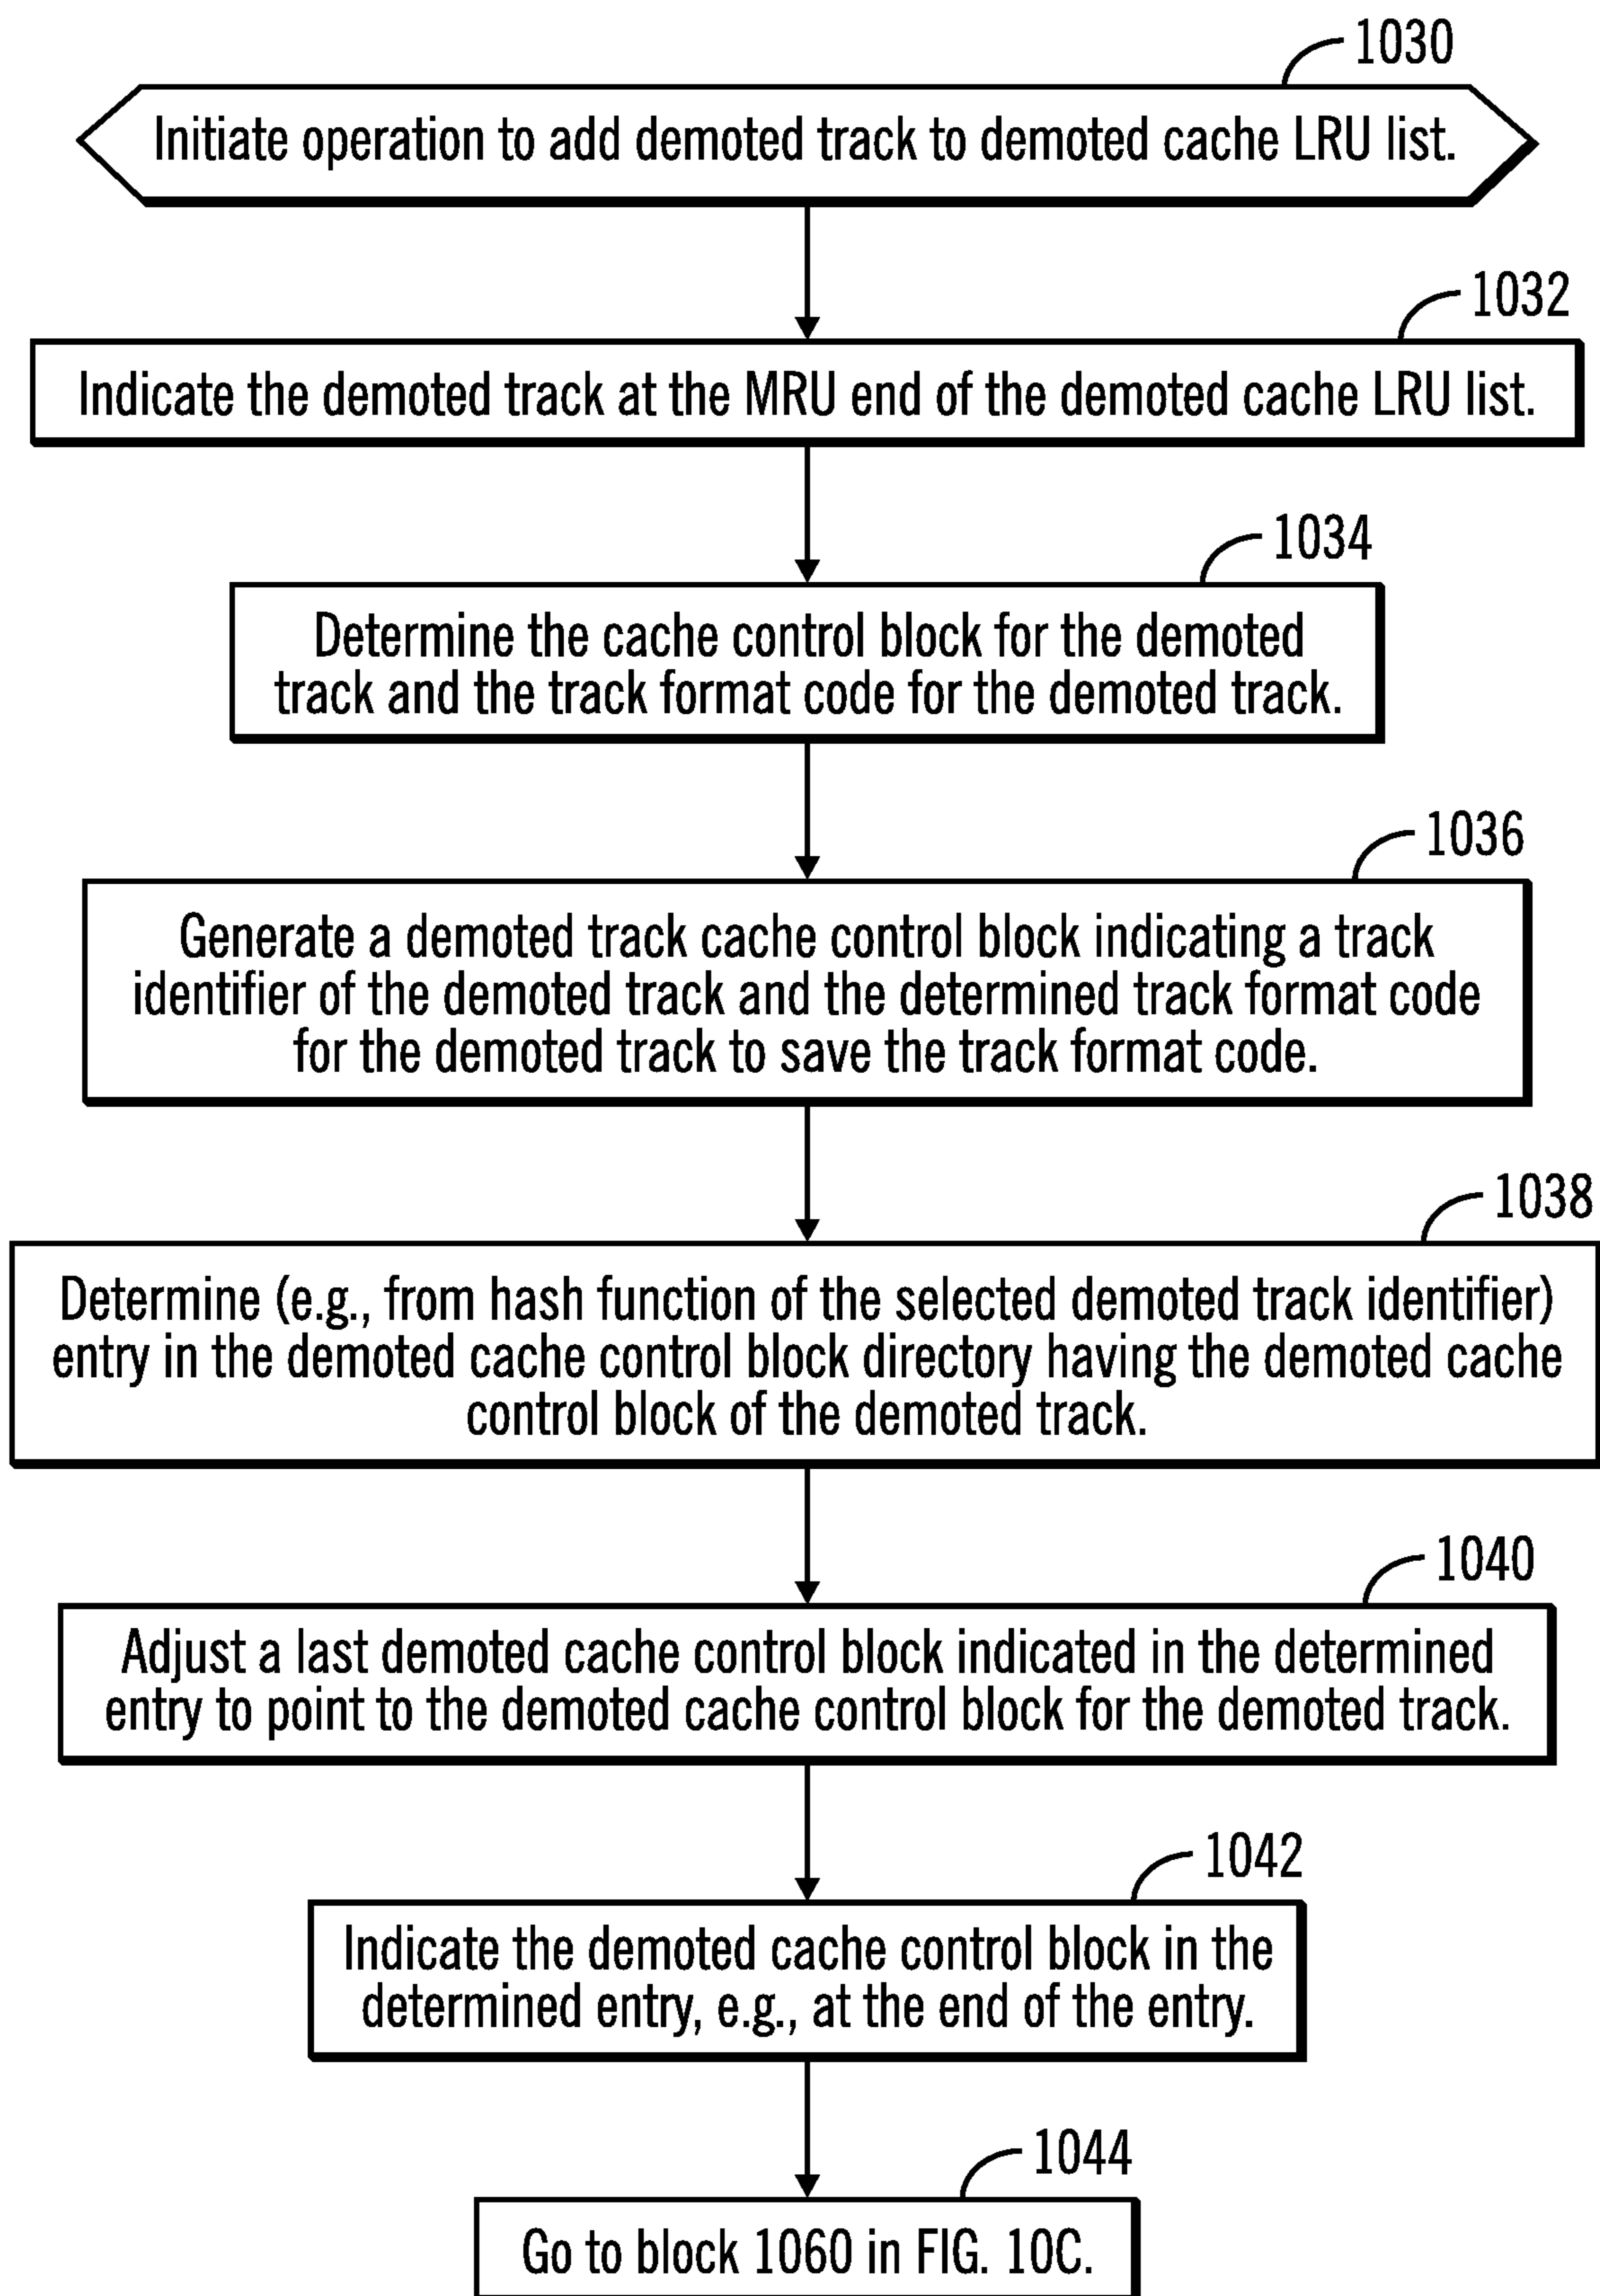
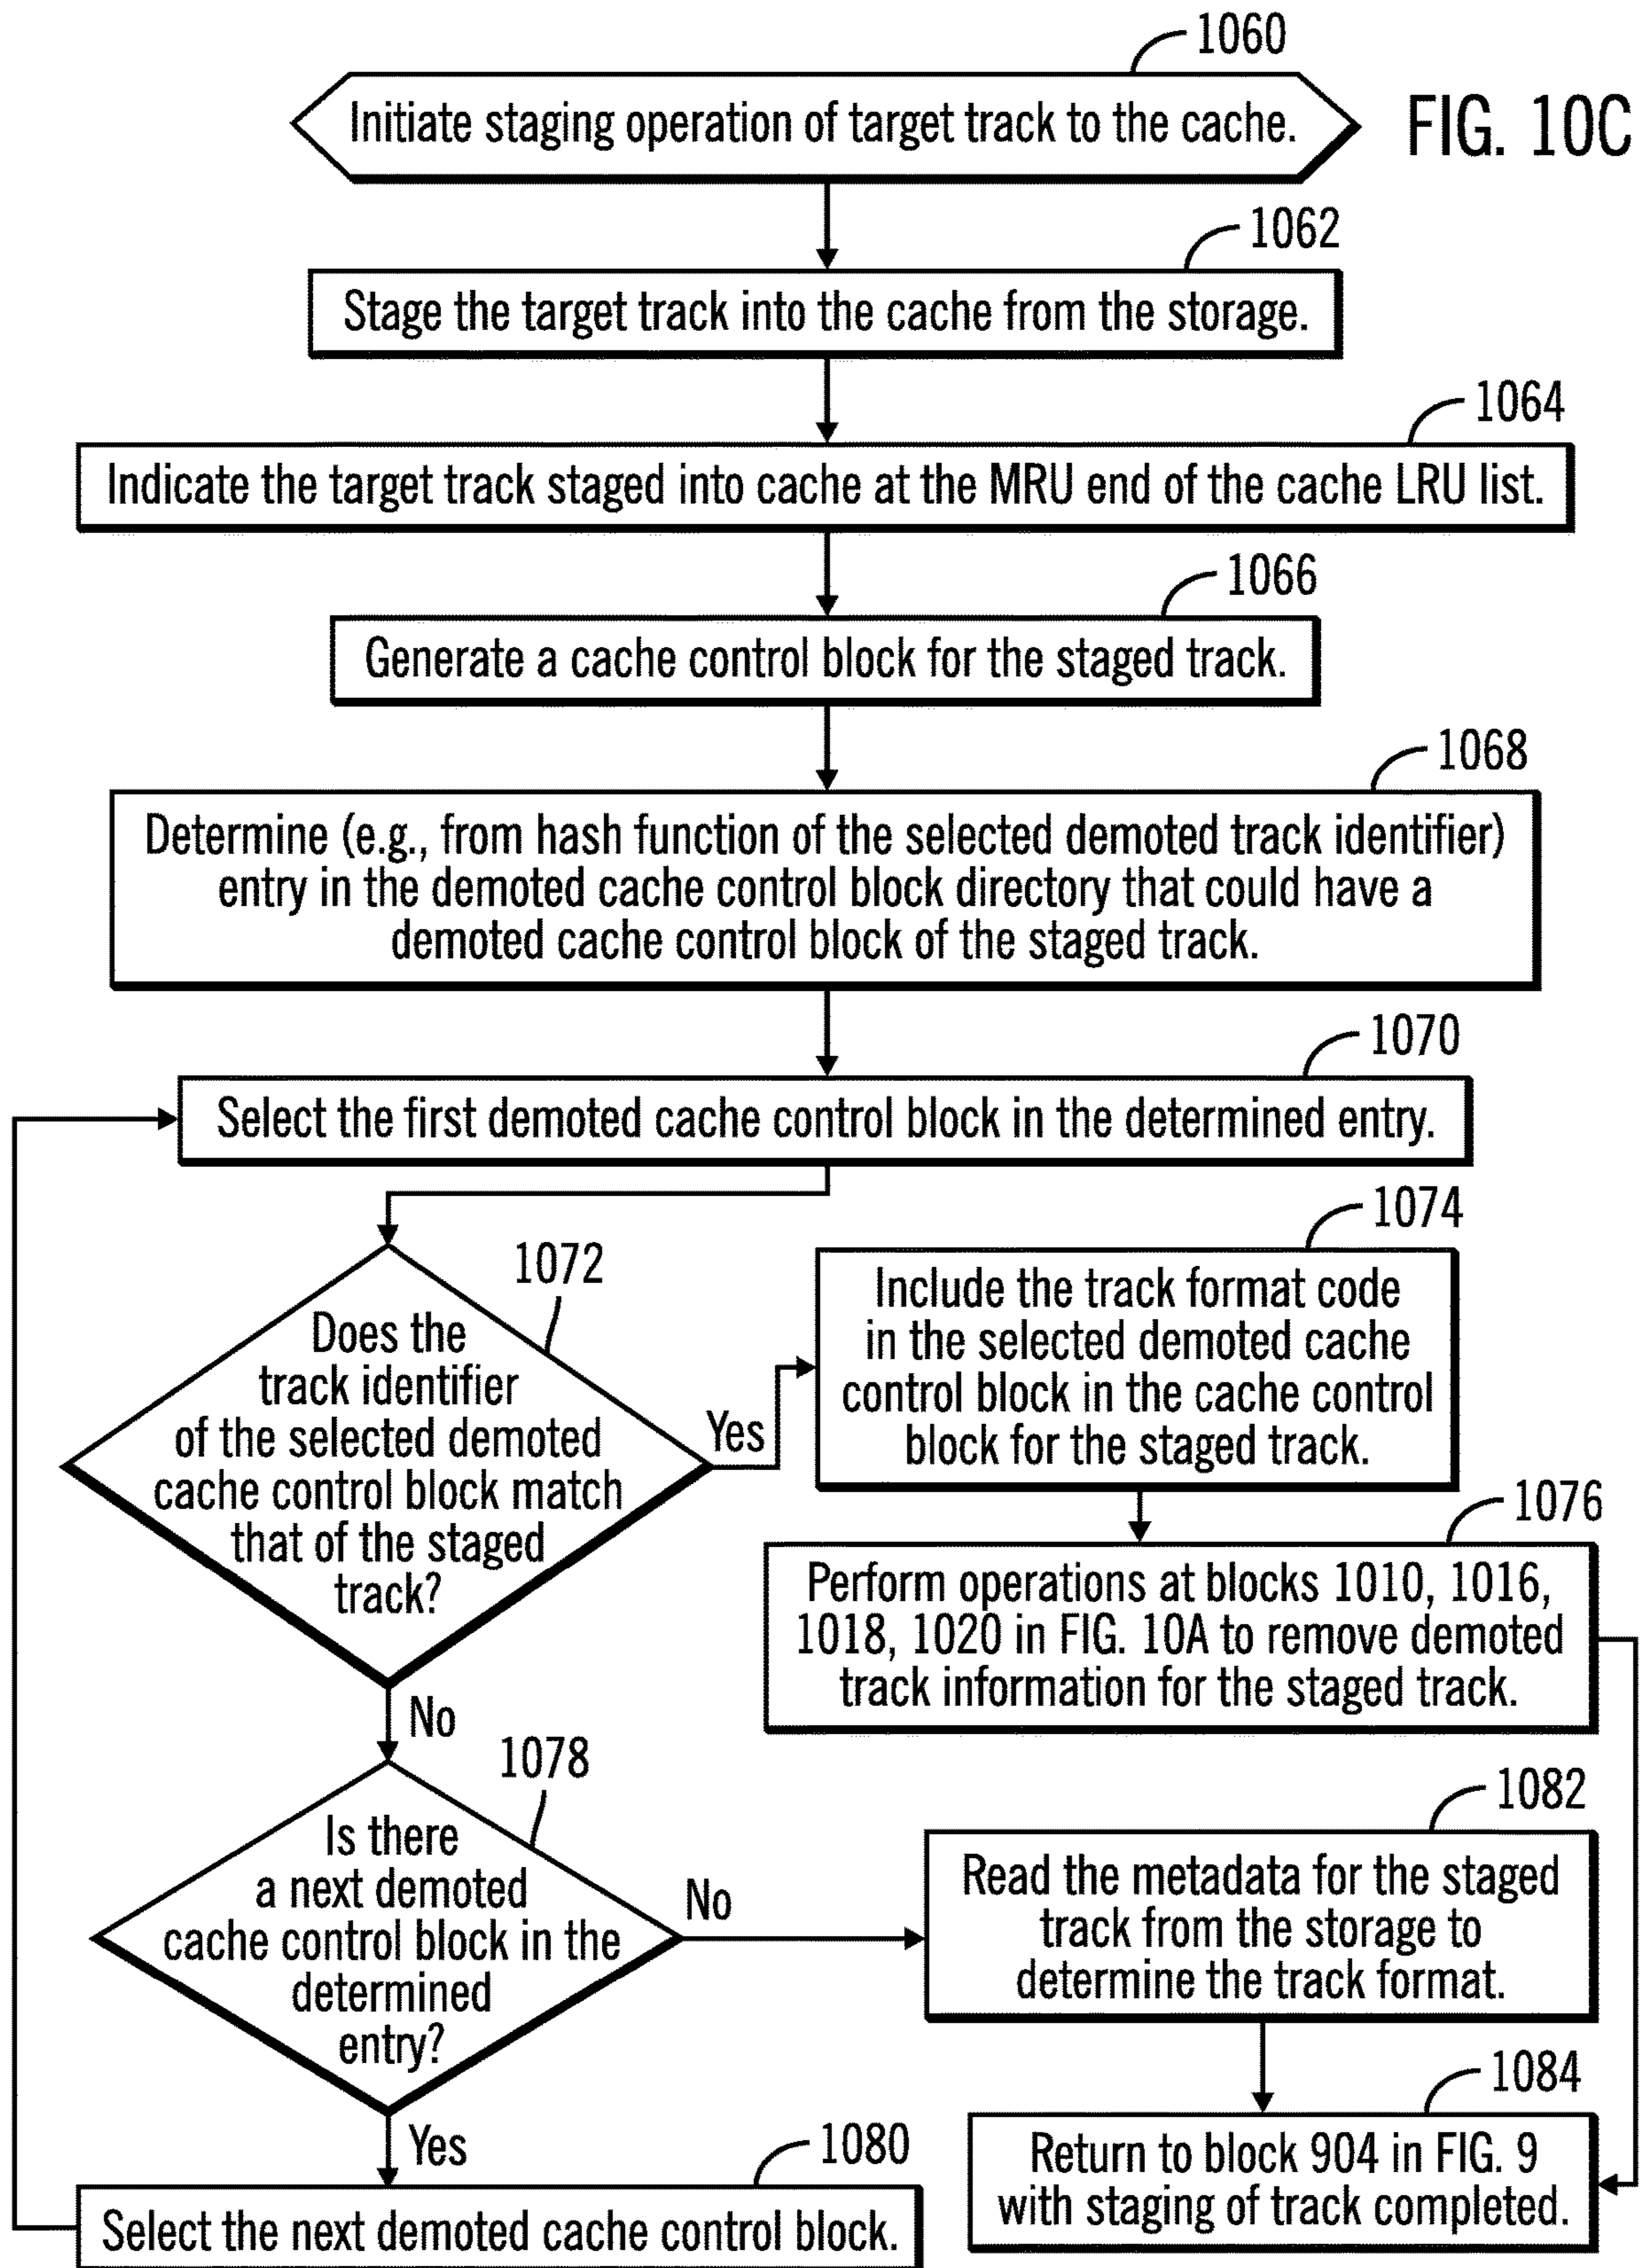


FIG. 10B



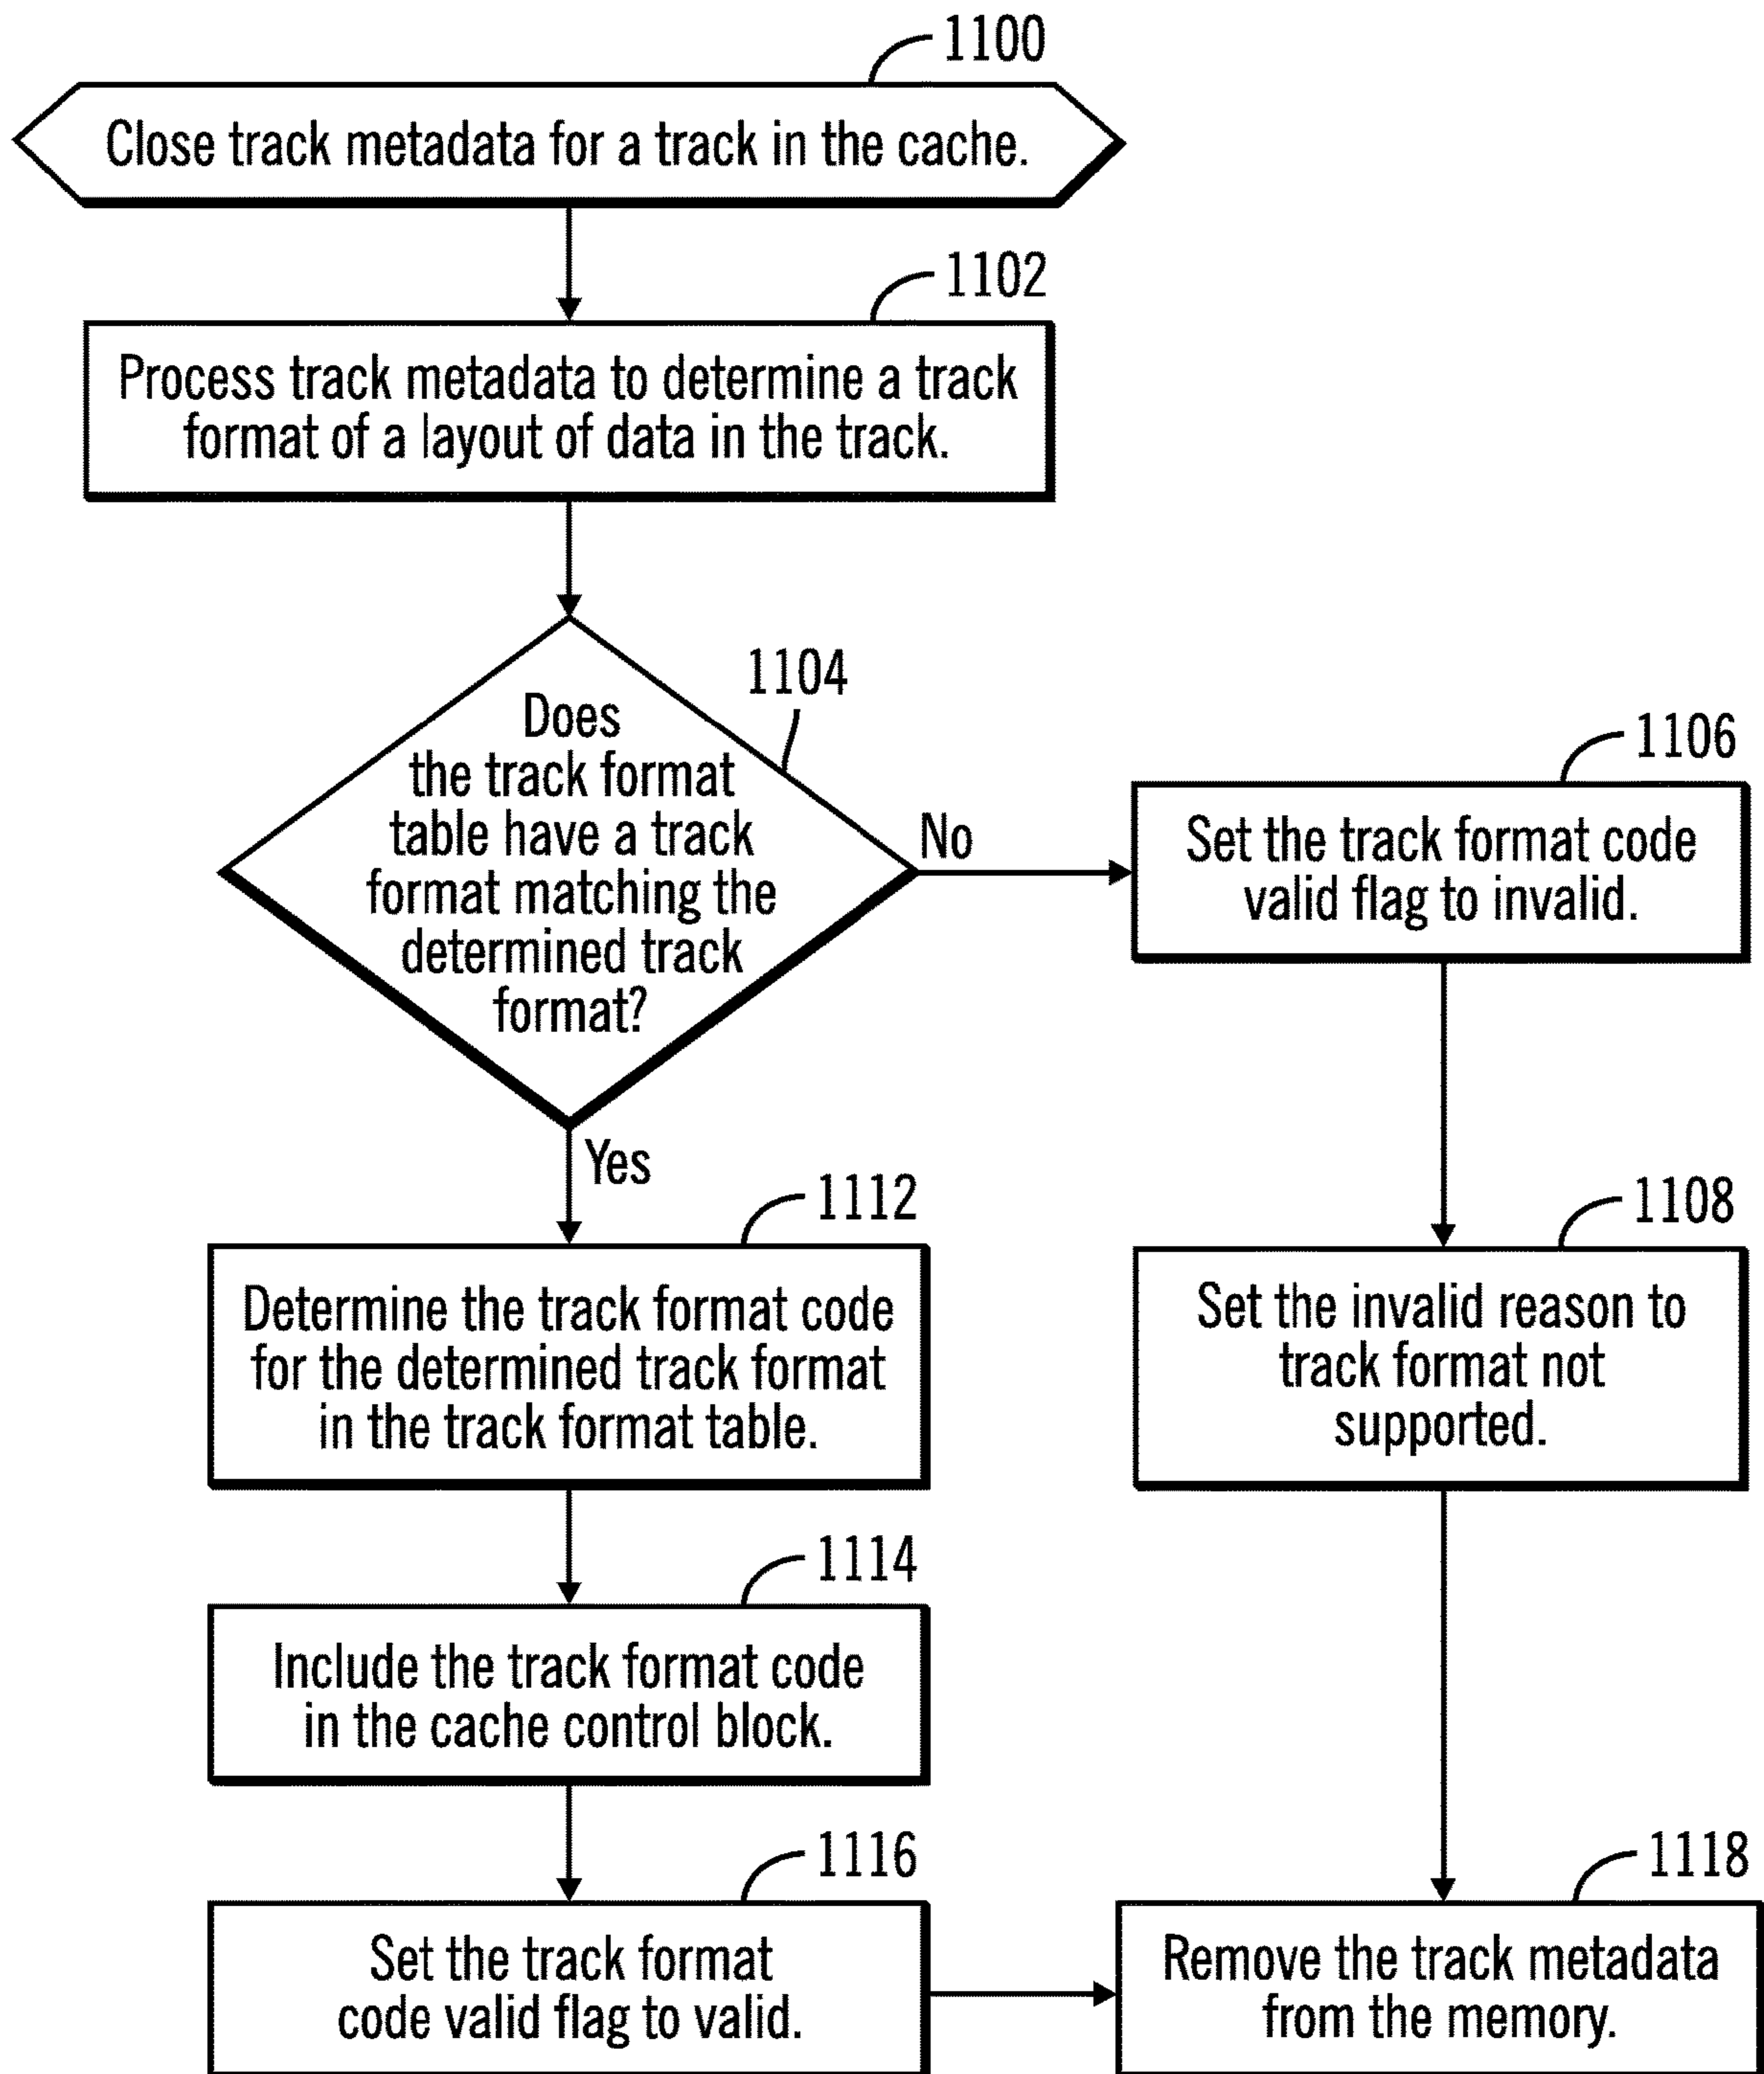


FIG. 11

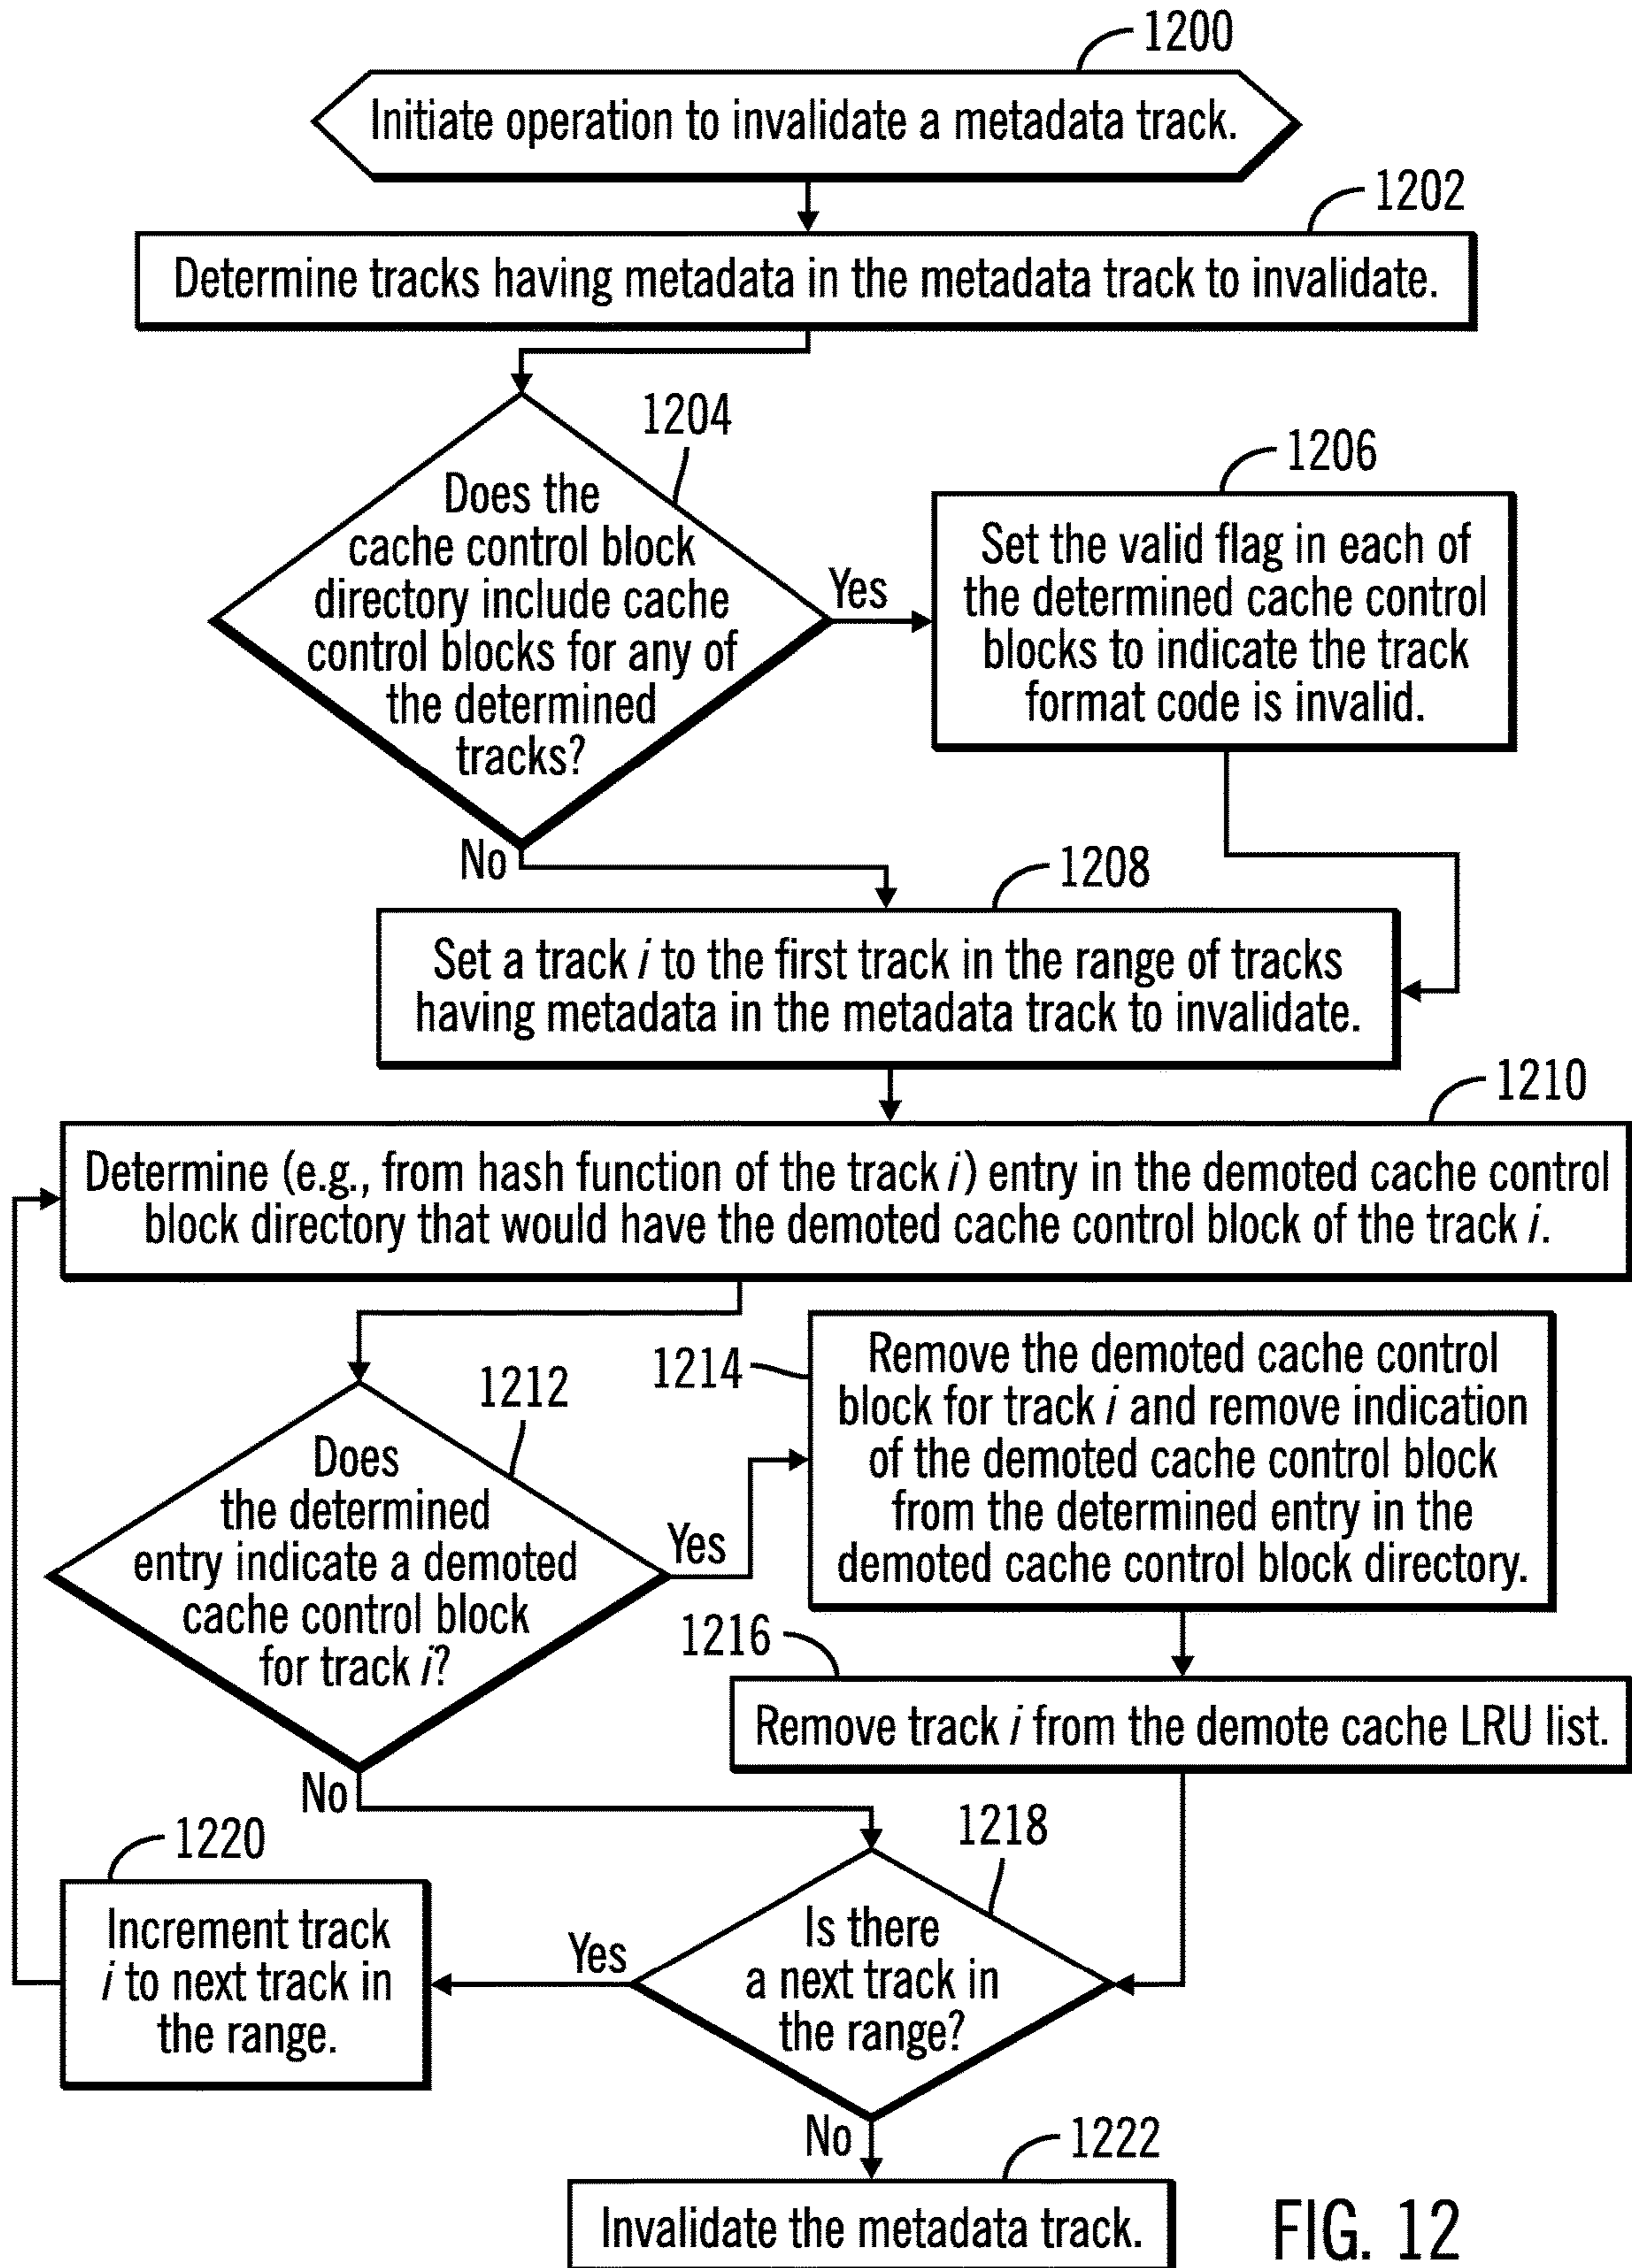


FIG. 12

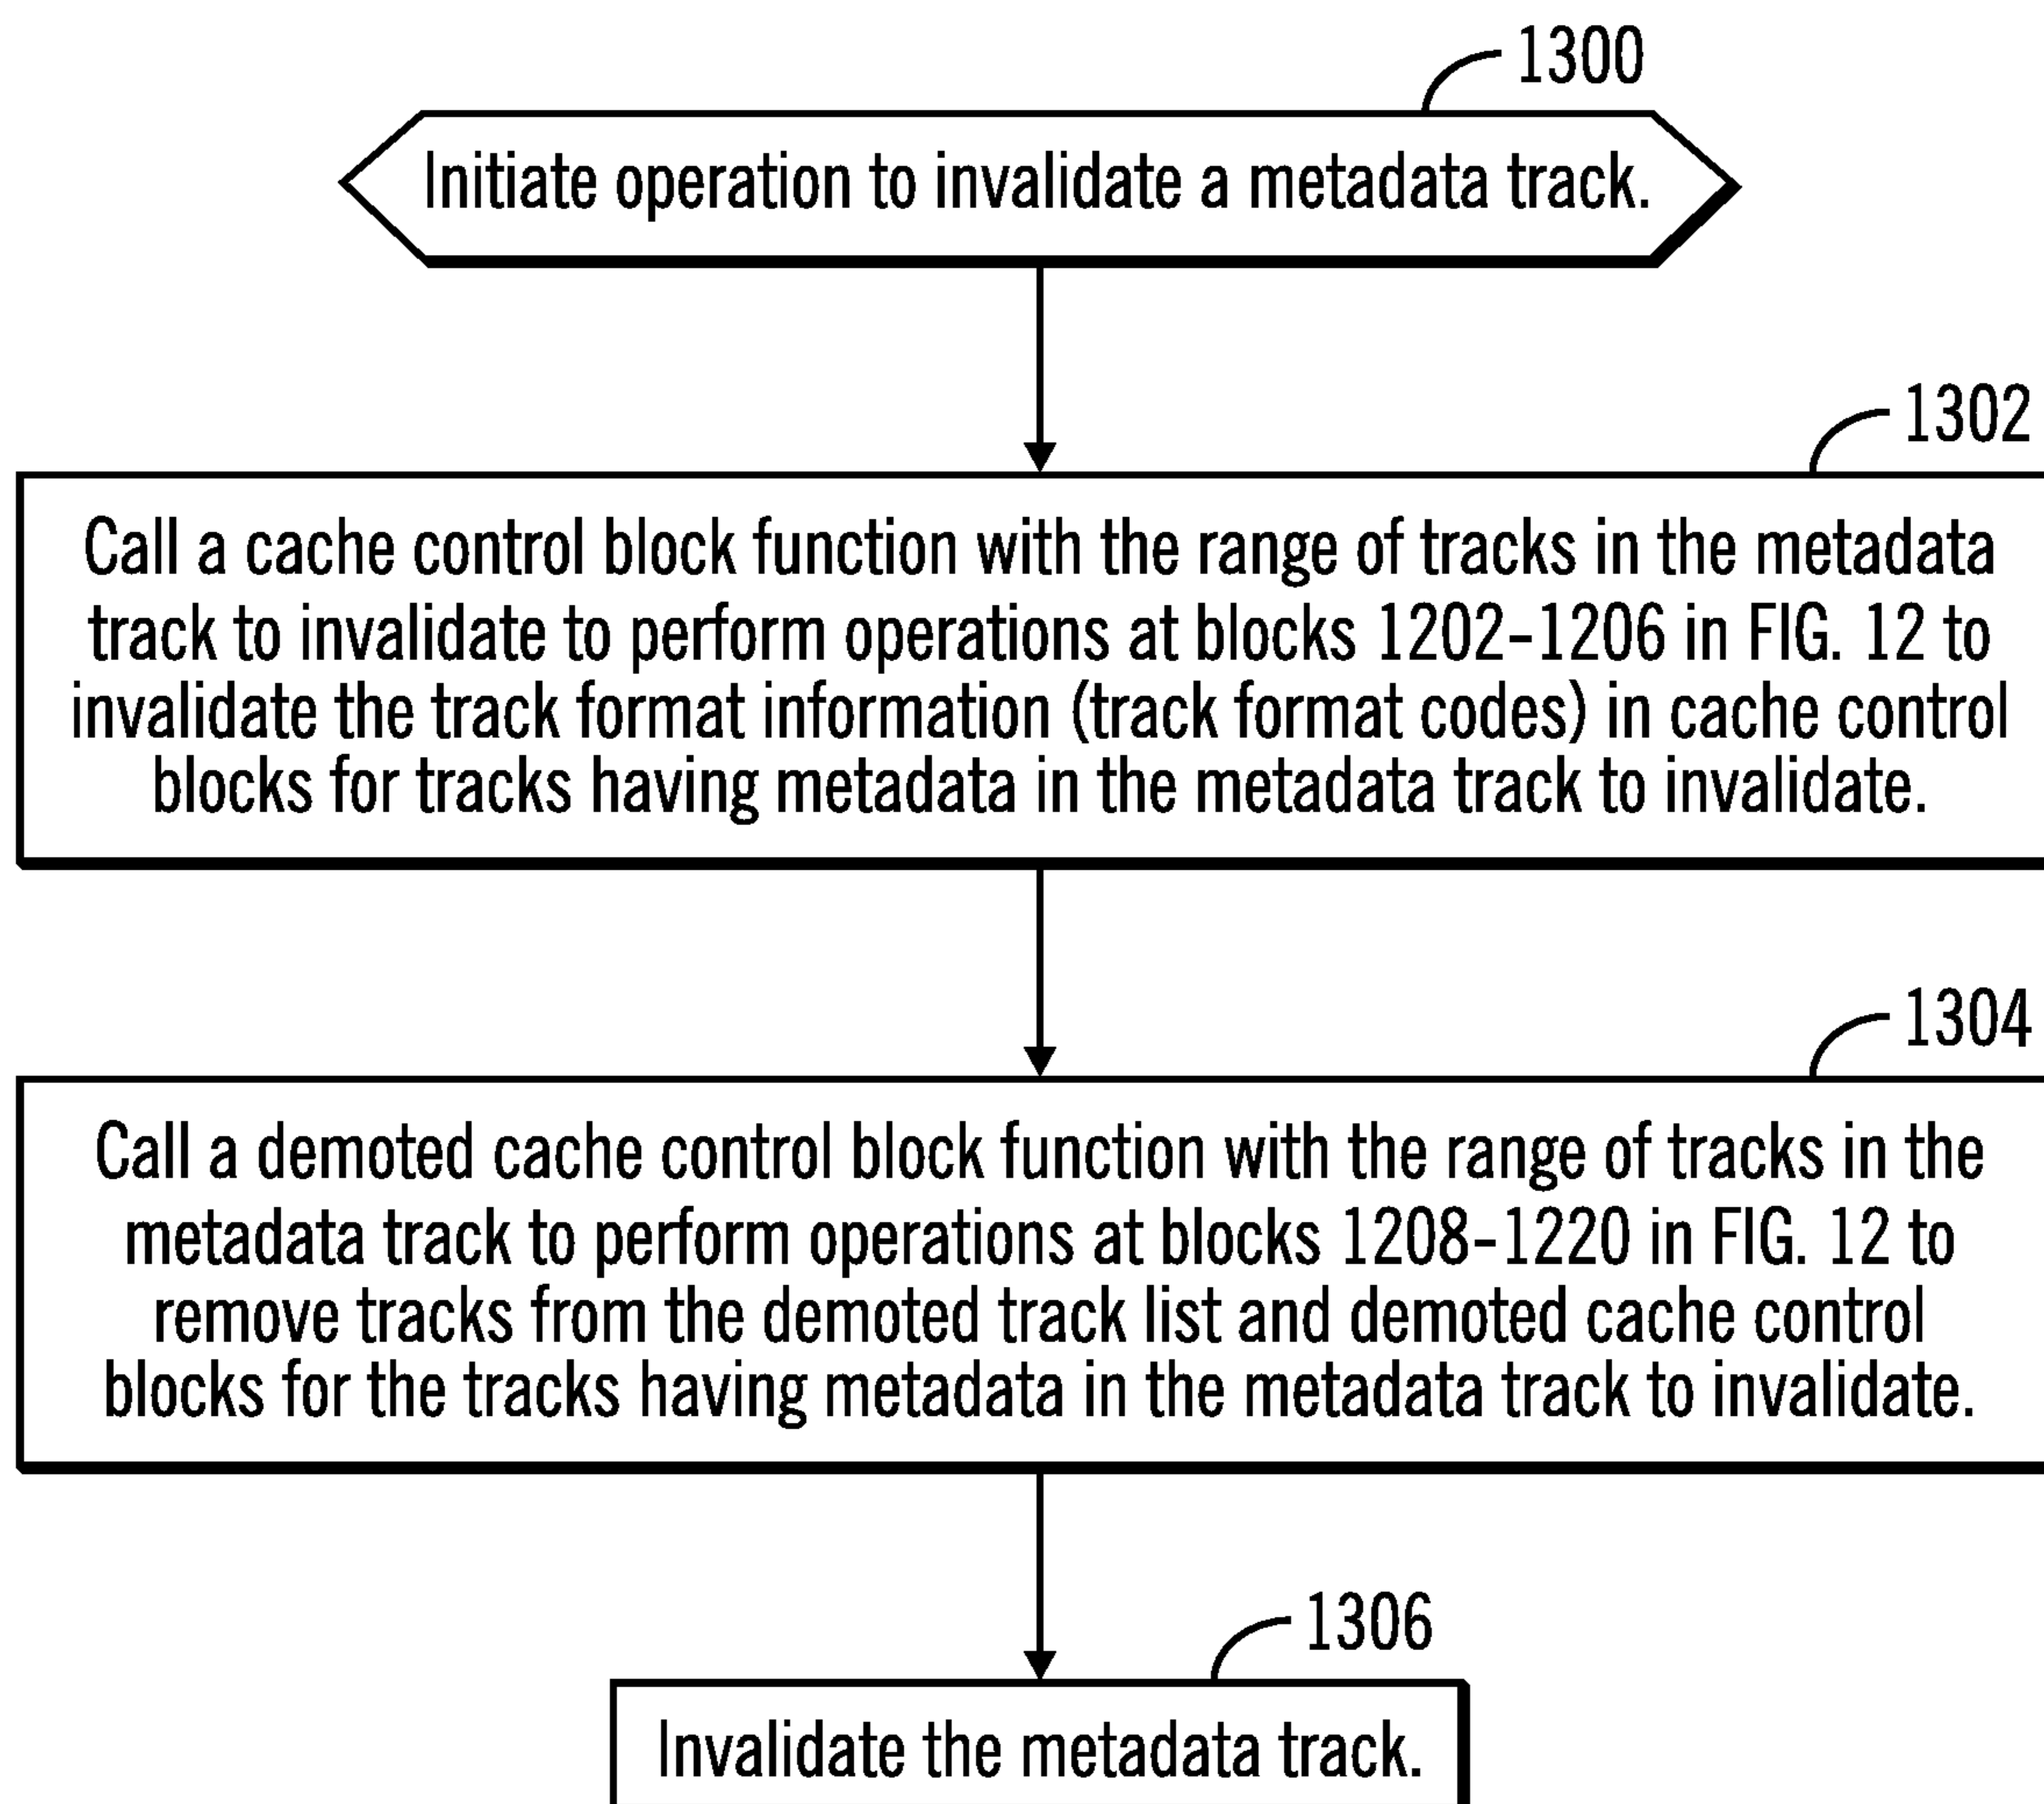


FIG. 13

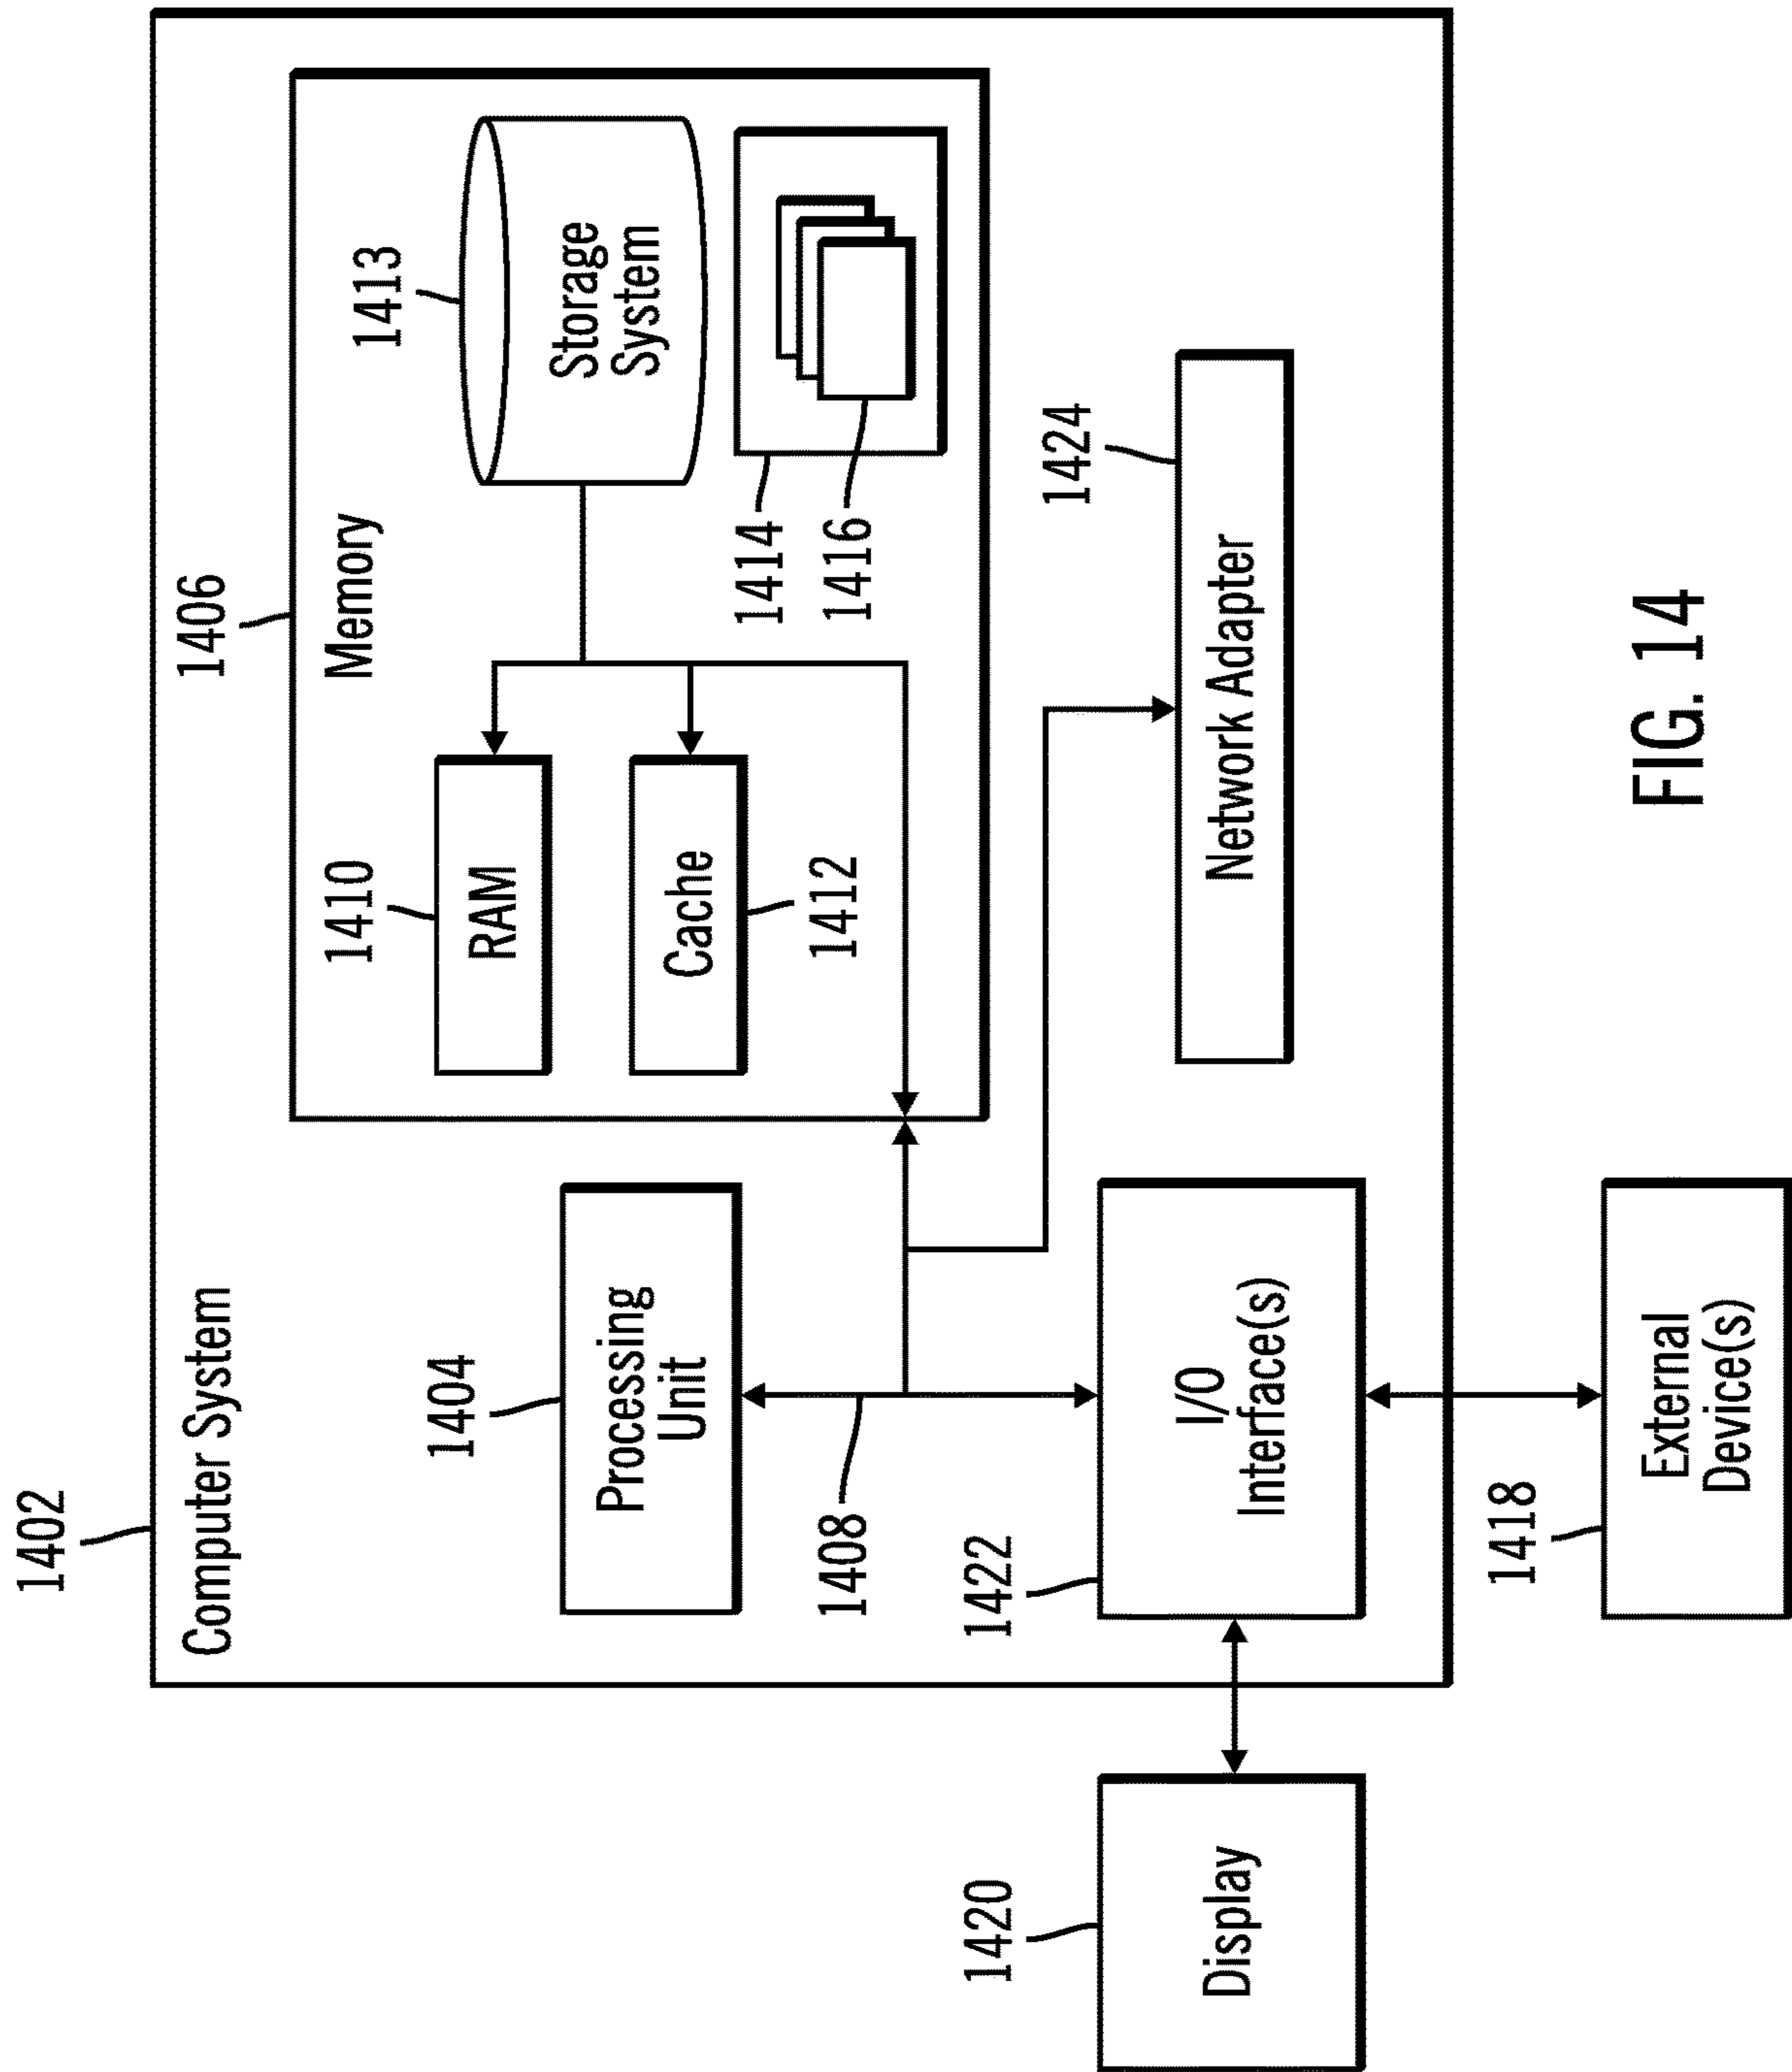


FIG. 14

1

**INVALIDATING TRACK FORMAT
INFORMATION FOR TRACKS IN CACHE**

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to a computer program product, system, and method for invalidating track format information for tracks in cache.

2. Description of the Related Art

In a storage environment, a host system may communicate a read/write request to a connected storage system over network channel through a network adaptor. If the data is in a cache of the storage system, i.e., a read hit, then the data may be returned quickly to the host system. This reduces the delay in returning requested data to a host I/O request. However, if the requested data is not in the cache of the storage system, then there may be significant latency realized while the storage system needs to retrieve the requested data from storage to return. Further, the thread or task executing the host read request may have to be context switched and deactivated in order to allow the host system to process further I/O requests. When the data is returned to the read request, then the task must be reactivated and data for the task must be returned to registers and processor cache to allow processing of the returned data for the read request.

The storage system stores track metadata for every track in cache that indicates a track format and layout of the track, also known as a track format descriptor (TFD), that is needed to process the data in the tracks in the cache. When processing a track in the cache, the track metadata needs to be read from storage and processed to determine the track format and layout to use to process requests to the track in the cache. A metadata track in storage may store metadata for numerous consecutive tracks in the storage. When the metadata track is invalidated, the metadata for all the tracks represented in the metadata track is invalidated. Further, the metadata track may maintain a flag for each track represented in the metadata track to allow for separately indicating the metadata for each track represented in the metadata track as invalid or valid.

There is a need in the art for improved techniques for managing and invalidating metadata in metadata tracks.

SUMMARY

Provided are a computer program product, system, and method for invalidating track format information for tracks in cache. Demoted tracks demoted from the cache are indicated in a demoted track list. Track format information is saved for the demoted tracks. The track format information indicates a layout of data in the demoted tracks, wherein the track format information for the demoted tracks is used when the demoted tracks are staged back into the cache. An operation is initiated to invalidate a metadata track of the metadata tracks in the storage. Demoted tracks indicated in the demoted track list having metadata in the metadata track to invalidate are removed. The track format information for the demoted tracks having metadata in the metadata track to invalidate is removed.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an embodiment of a storage environment.

2

FIG. 2 illustrates an embodiment of a track format table entry.

FIG. 3 illustrates an embodiment of a cache control block.

FIG. 4 illustrates an embodiment of a cache Least Recently Used (LRU) list of tracks in the cache.

FIG. 5 illustrates an embodiment of a demoted cache Least Recently Used (LRU) list of tracks demoted from the cache.

FIG. 6 illustrates an embodiment of a demoted cache control block.

FIG. 7 illustrates an embodiment of a demoted cache control block directory entry.

FIG. 8 illustrates an embodiment of operations to process a read/write request received on a first channel, such as a bus interface.

FIG. 9 illustrates an embodiment of operations to process a read/write request received on a second channel, such as a network.

FIGS. 10a, 10b, and 10c illustrate an embodiment of operations to stage a track into the cache.

FIG. 11 illustrates an embodiment of operations to close track metadata and determine a track format code for the track in cache of the closed track metadata.

FIGS. 12 and 13 illustrate embodiments of operations to invalidate track format information for tracks having metadata in a metadata track to invalidate.

FIG. 14 illustrates an embodiment of a computer architecture used with described embodiments.

DETAILED DESCRIPTION

In a storage environment, a host system may first communicate a read/write request to a connected storage system over a fast channel, such as a bus interface, such as the Peripheral Component Interconnect Express (PCIe) interface. For a read/write request over the fast channel which is supposed to complete within a threshold time, the host system holds the application thread for the read/write request in a spin loop waiting for the request to complete. This saves processor time associated with a context swap, which deactivates the thread and reactivates the thread in response to an interrupt when a response to the read/write request is received. If the data for the read/write request sent on the fast channel is not in cache, then the storage system may fail the read/write request and the host system may communicate the same read/write request over a storage area network via a host adaptor, which is slower than processing the I/O request over the PCIe interface. Communicating the read/write request over the second channel requires the host system to perform context switch from the task handling the read/write request to another task while waiting for the read/write request to complete. Context switching is costly because it requires the processor running the task to clear all registers and L1 and L2 caches for the new task, and then when completing the new task, reactivate the context switched task and return the state data to the registers and L1 and L2 caches for the task that was context switched while waiting for the read/write request to complete.

Certain read/write operations need to be completed within a threshold time, else they are failed. The storage system will have to access track metadata to process a request to a track. The track metadata provides information on the format of data and layout of records in the track that are needed in order to perform reads and writes to the track. However, the reading of the track metadata from the storage comprises a substantial portion of the latency in processing read/write request. Described embodiments provide improvements to

cache technology by reducing cache operation latency by including a track format code in the cache control block for a track in the cache. This track format code may be used for fast access to the track format from a track format table without having to read the track metadata from storage. By 5 eliminating the need to read the track metadata to determine the track layout, described embodiments increase the likelihood that read/write requests on the first channel that need to be completed within a threshold time are completed by accessing the track metadata for a track from the track 10 format table, associating track format codes with common track formats.

Described embodiments provide further improvements to cache computer technology by reducing cache latency for a track staged into cache that was previously demoted by 15 saving track format metadata, such as the track format code, when a track is demoted from the cache. When the demoted track is later staged into cache, the track format metadata may be quickly determined by the track format information saved with the demoted track and included in a cache control 20 block for the staged track without having to read the metadata for the staged track. Avoiding the need to read the track metadata for a staged track substantially reduces the latency and delays in staging a track and processing a read/write request to the staged track.

With described embodiments, a read/write request to a target track on a channel requiring that the request be completed within a threshold time is processed if the track format code for the target track is within the cache control 25 block for the target track. Using the track format code to access the track format from the track format table reduces the latency of access to the track format metadata to allow the read/write request to complete within the time threshold. This keeps the time the host thread is spinning on the read/write request task within an acceptable time threshold. However, if the cache control block for the target track does not have a valid track format code, then the read/write request on the first channel is failed because it is unlikely the read/write request can complete within the threshold time 30 given that the track format metadata will have to be retrieved from the storage. Failing the read/write request on the first channel, causes the host to re-drive the read/write request on the second channel. The processing of the read/write request on the second channel reads in the track metadata from the storage to allow for processing the read/write request and adding the track format code to the cache control block for the target track.

Described embodiments further provide improvements to computer technology to remove track format information maintained for tracks that are in cache or have been recently demoted from cache when invalidating a metadata track 35 having the metadata for the tracks. Metadata tracks may be invalidated when track space is being released or when the metadata information is no longer valid due to a change in the track layout or format in storage. When track format information is maintained for tracks in the cache or recently demoted from cache in cache control blocks and demoted cache control blocks, then that track format information needs to be invalidated and made unavailable since that track format information will no longer be valid after the 40 metadata track is invalidated. For this reason, described embodiments provide techniques for invalidating all instances of track format information maintained for tracks in cache and for tracks demoted from cache before the metadata track is invalidated, so that invalid track format information will not remain in the system and be referenced for track access after the metadata track is invalidated.

FIG. 1 illustrates an embodiment of a storage environment in which a host 100 directs read and write requests to a storage system 102 to access tracks in volumes configured in storage devices 104 in a disk enclosure 106. The host 100 includes a processor complex 108 of one or more processor 5 devices and a memory 110 including an operating system 111 executed by the processor complex 108. The host operating system 111 generates read and write requests to tracks configured in the storage devices 104. The host 100 includes hardware to communicate read and write requests on two different channels. A first channel is a bus interface, such as a Peripheral Component Interconnect Express (PCIe), including a bus 112, a bus switch 114 to connect one or more devices on the bus 112, including the processor 10 complex 108, a memory system 110, and a bus host adaptor 116 to extend the bus interface over an external bus interface cable 118 to the storage system 102. A second channel to connect the host 100 and storage system 102 uses a network 10 host adaptor 120, connected to the bus 112, that connects to a separate network 122 over which the host 100 and storage 20 system 102 additionally communicate. The first channel through the bus interface may comprise a faster access channel than the network 122 interface through the network host adaptor 120.

The storage system 102 includes a bus interface comprising a bus 124a, 124b, a bus switch 126 to connect to endpoint devices on the bus 124a, 124b, and a bus host adaptor 128 to connect to the external bus interface cable 118 to allow communication over the bus interface to the 25 host 100 over the first channel. The storage system 102 includes an Input/Output bay 130 having the bus host adaptor 128, one or more device adaptors 132 to connect to the storage devices 104, and one or more network host adaptors 134 to connect to the network 122 and host systems.

The storage system 102 includes a processor complex 136 of one or more processor devices and a memory 138 having a cache 140 to cache tracks accessed by the connected hosts 100. The memory 138 includes a cache manager 142 and a storage manager 144. The storage manager 144 manages access requests from processes in the hosts 100 and storage 30 system 102 for tracks in the storage 104. The devices 136, 138, 128, 132, and 134 connect over the bus interface implemented in the bus lanes 124a, 124b and bus switch 126.

The cache manager 142 maintains accessed tracks in the cache 140 for future read access to the tracks to allow the accessed tracks to be returned from the faster access cache 140 instead of having to retrieve from the storage 104. Further, tracks in the cache 140 may be updated by writes. A track may comprise any unit of data configured in the storage 104, such as a track, Logical Block Address (LBA), etc., which is part of a larger grouping of tracks, such as a volume, logical device, etc.

The cache manager 142 maintains cache management information 146 in the memory 138 to manage read (unmodified) and write (modified) tracks in the cache 140. The cache management information 146 may include a track format table 200 having track format codes for common track format descriptors providing details of a layout and format of data in a track; track index 148 providing an index of tracks in the cache 140 to cache control blocks in a control block directory 300; and a cache Least Recently Used (LRU) list 400 for tracks in the cache 140. The control block 35 directory 300 includes the cache control blocks, where there is one cache control block for each track in the cache 140 providing metadata on the track in the cache 140. The track

5

index **148** associates tracks with the cache control blocks providing information on the tracks in the cache **140**. Upon determining that the cache LRU list **400** is full or has reached a threshold level, tracks are demoted from the LRU list **400** to make room for more tracks to stage into the cache **140** from the storage **104**.

In certain embodiments, there may be multiple hosts **100** that connect to the storage system **102** over the first and second channels to access tracks in the storage devices **104**. In such case, the storage system **102** would have at least one bus host adaptor **128** to connect to the bus interface **118** of each connected host **100** and one or more network host adaptors **134** to connect to the network host adaptors **120** on the hosts **100**.

In one embodiment, the bus interfaces **112**, **114**, **116**, **118**, **124a**, **124b**, **126**, and **128** may comprise a Peripheral Component Interconnect Express (PCIe) bus interface technology. In alternative embodiments, the bus interfaces **112**, **114**, **116**, **118**, **124a**, **124b**, **126**, and **128** may utilize suitable bus interface technology other than PCIe. The bus host adaptors **116** and **128** may comprise PCIe host adaptors that provide the interface to connect to the PCIe cable **118**. The network **122** may comprise a Storage Area Network (SAN), a Local Area Network (LAN), a Wide Area Network (WAN), the Internet, an Intranet, etc., and the network host adaptors **120**, **134** provide the network **122** connections between the hosts **100** and storage system **102**.

The storage system **102** may comprise a storage system, such as the International Business Machines Corporation (IBM®) DS8000® and DS8880 storage systems, or storage controllers and storage systems from other vendors. (IBM and DS8000 are trademarks of International Business Machines Corporation throughout the world). The host operating system **111** may comprise an operating system such as Z Systems Operating System (Z/OS®) from International Business Machines Corporation (“IBM”) or other operating systems known in the art. (Z/OS is a registered trademark of IBM throughout the world).

The storage devices **104** in the disk enclosure **106** may comprise different types or classes of storage devices, such as magnetic hard disk drives, solid state storage device (SSD) comprised of solid state electronics, EEPROM (Electrically Erasable Programmable Read-Only Memory), flash memory, flash disk, Random Access Memory (RAM) drive, storage-class memory (SCM), etc., Phase Change Memory (PCM), resistive random access memory (RRAM), spin transfer torque memory (STT-RAM), conductive bridging RAM (CBRAM), magnetic hard disk drive, optical disk, tape, etc. Volumes in a storage space may further be configured from an array of devices, such as Just a Bunch of Disks (JBOD), Direct Access Storage Device (DASD), Redundant Array of Independent Disks (RAID) array, virtualization device, etc. Further, the storage devices **104** in the disk enclosure **106** may comprise heterogeneous storage devices from different vendors and different types of storage devices, such as a first type of storage devices, e.g., hard disk drives, that have a slower data transfer rate than a second type of storage devices, e.g., SSDs.

FIG. 2 illustrates an embodiment of a track format table entry **200_i** in the track format table **200**, which includes a track format code **202** and the track format metadata **204**. In certain embodiments Count Key Data (CKD) track embodiments, the track format metadata **204** may comprise a track format descriptor (TFD) indicating a number of records in the track, a block size, a number of blocks in the track, a data length of each of the records, and a control interval size indicating an amount of data that is read or written atomi-

6

cally as a unit, number of blocks in a control interval, and whether a control interval spans two tracks, and other information. The track format code **202** may comprise an index value of the index entry **200_i** in the track format table **200**. For instance, if there are 32 track format table entries **200_i**, then the track format code **202** may comprise 5 bits to reference the different possible number of 32 entries **200_i**.

FIG. 3 illustrates an embodiment of a cache control block **300_i** for one of the tracks in the cache **140**, including, but not limited to, a cache control block identifier **302**, such as an index value of the cache control block **300_i**; a track ID **304** of the track in the storage **104**; the cache LRU list **306** in which the cache control block **300_i** is indicated; an LRU list entry **308** at which the track is indicated; a cache timestamp **310** indicating a time the track was added to the cache **140** and indicated on the LRU list **306**; additional track metadata **312** typically maintained for tracks stored in the cache **140**, such as a dirty flag indicating whether the track was modified; a track format code **314** comprising one of the track format codes **202** of the track format metadata **204** describing the layout of data in the track **304** represented by the cache control block **300_i**; a track format code valid flag **316** indicating whether the track format code **314** is valid or invalid; and an invalid reason **318** indicating a reason for the track format code valid flag **316** code being invalid, as indicated in the track format code valid flag **316**.

FIG. 4 illustrates an embodiment of the cache LRU list **400**, having a most recently used (MRU) end **402** identifying a track most recently added to the cache **140** or most recently accessed in the cache **140** and a least recently used (LRU) end **404** from which the track identified at the LRU end **404** is selected to demote from the cache **140**. The LRU end **404** points to a track identifier, such as a track identifier address or a cache control block **300_i** for the track, of the track that has been in the cache **140** the longest for tracks indicated in that list **400**.

FIG. 5 illustrates an embodiment of the demoted cache LRU list **500**, having a most recently used (MRU) end **502** identifying a demoted track most recently added to the demoted cache LRU list **500** and a least recently used (LRU) end **504** from which the demoted track identified at the LRU end **504** is selected to demote from the demoted cache LRU list **500**.

FIG. 6 illustrates an embodiment of a demoted cache control block **600_i** having a track identifier (ID) **602** of a demoted track; a track format code **604** if available of the track format metadata **204** in the track format table **200** of the demoted track; a pointer to a previous LRU entry **606** of a previous demoted track in the demoted cache LRU list **500**; a pointer to a next LRU entry **608** of a next demoted track in the demoted cache LRU list **500**; and a pointer **610** to next demoted cache control block in the entry in the demoted cache control block directory **700** for the track ID **602**.

In additional embodiments, the track format code **604** may comprise track format information other than a code **604** in a table **200**, such as other information that may be used to identify or represent the track format metadata and layout of data in the tracks and comprises substantially less bits of information than the represented track format metadata and layout information.

FIG. 7 illustrates an embodiment of an entry **700_i** in the demoted cache control block **700** that includes pointers **702₁**, **702₂** . . . **702_n** to demoted cache control blocks **600_i**. Each demoted cache control block **600_i** maps to one entry in the demoted cache control block directory **700** based on the track ID **602**. In one embodiment, a hash function of the

track identifier would produce the entry in the demoted cache control block directory 700 in which the demoted cache control block 600_i is indicated.

FIG. 8 illustrates an embodiment of operations performed by the cache manager 142 and storage manager 144 to process a read/write request to a target track received on a first fast channel, such as the PCIe bus interface via bus host adaptor 128. Upon receiving (at block 800) the read/write request at the bus host adaptor 128, if (at block 802) the target track is not in the cache 140, then the storage manager 144 returns (at block 804) fail to the read/write request on the first channel or bus host adaptor 128 to the host 100, which causes the host 100 to retry the read/write request on the second channel or network host adaptor 120, 134. Failure is returned because if the target track is not in the cache 140, then the target track and track metadata needs to be staged into cache 140, which would exceed the time threshold for read/writes on the first channel, where the host processor is spinning on the thread of the read/write request. If (at block 802) the target track is in the cache 140 is a write and if (at block 808) the write modifies the track format, then the cache manager 142 sets (at block 810) the track format code valid flag 316 to invalid and indicates (at block 812) the invalid reason 318 that the track in the cache 140 was invalidated. The storage manager 144 then returns (at block 804) fail to the host 100 because the track metadata needs to be read from the storage 104 to update with the modified track format.

If (at block 806) the read/write request is a read or if (at block 808) the request is a write that does not modify the track format, then the cache manager 142 determines (at block 814) if the track format code valid flag 316 is set to valid. If so, then the cache manager 142 determines (at block 816) the track format metadata 204 in the track format table 200 corresponding to the track format code 314 in the cache control block 300_i. The cache manager 142 uses (at block 818) the track format layout indicated in the determined track format metadata 204 to process the read or write request to the target track in the cache 140. If the request is a write, a dirty flag 312 in the cache control block 300_i may be set to indicate the track is modified. If (at block 814) the track format code valid flag 316 is invalid, meaning there is no fast access to track metadata available through the track format code 314, then the storage manager 144 returns (at block 804) fail on the bus interface to the bus host adaptor 128 because the track format table 200 cannot be used, and the track metadata needs to be read from the storage 104, which would introduce too much latency for the fast read/write on the first channel.

With the embodiment of operations of FIG. 8, during a fast write over the bus interface or first channel, if the track format metadata may be accessed without latency through the track format table 200, then the read/write request is allowed to proceed because the transaction can be processed very quickly because the track metadata can be obtained directly from the track format table 200 through the track format code 314, without having to be read. However, if the cache control block 300_i does not have a valid track format code 314 to allow low latency access of track metadata, then the read/write request is failed because the transaction will not likely complete within a fast time threshold. This determination is important to avoid host delays in processing other tasks while the host processor is spinning on the thread handling the read/write request while waiting for the read/write request to complete. If the track metadata can be accessed from the track format table 200 than there is a high likelihood the read/write can complete on the bus interface

channel within the time required to avoid the host processor holding the thread for too long, which causes other I/O requests to be queued and delayed. If the track metadata cannot be accessed from the track format table 200 and needs to be read from the storage 104, then it is unlikely the read/write request will complete within the time threshold for the host processor to spin on the thread for the read/write request, and failure is returned. Returning failure when the track metadata cannot be obtained from the track format table 200 causes the host thread waiting on the read/write request task to be deactivated and the host processor may context switch to processing other tasks, and then the read/write request is tried on the second network channel during the context switch.

FIG. 9 illustrates an embodiment of operations performed by the cache manager 142 and storage manager 144 to process a read/write request to a target track received on a second channel, such as the network 122 on network host adaptor 134. Upon receiving (at block 900) the read/write request, if (at block 902) the target track is not in the cache 140, then the cache manager 142 proceeds (at block 904) to block 1000 in FIG. 10a to stage the track into the cache 140. If (at block 908) the read/write request is a write and if (at block 910) the write modifies the track format, then the cache manager 142 updates (at block 912) the track metadata to indicate the modified track format and sets (at block 914) the track format code valid flag 316 to invalid. The track metadata 312 is further updated (at block 916) to indicate the track is modified or dirty. If (at block 908) the request is a read or from block 916, the cache manager 142 uses (at block 918) the track format layout indicated in the track format metadata to process the read or write request to the target track in the cache 140.

If (at block 902) the target track is in the cache 140 and if (at block 930) the track format code valid flag 316 is set to valid, then the cache manager 142 determines (at block 932) the track format metadata 204 in the track format table 200 corresponding to the track format code 314 in the cache control block 300_i for the target track. From block 932, control proceeds to block 908 to process the read/write request. If (at block 930) the track format code valid flag 316 is set to invalid, then the cache manager 142 reads (at block 934) the track metadata for the target track from the storage 104 to determine the track format, e.g., size of blocks, control interval, layout of records on the track, etc. From block 934, control proceeds to block 908 to process the read/write request.

With the embodiment of FIG. 9, when the read/write request is received on the second slower channel, such as over the network 122, where the host operating system 111 would have performed a context switch for the thread handling the read/write request, the cache manager 142 may read the track metadata from the storage 104 to determine the track layout to process the request. During this time, the host processing of further host requests is not delayed because the host thread handling the read/write request is context switched and not active, until the read/write request returns complete.

FIGS. 10a, 10b, and 10c illustrate an embodiment of operations performed by the cache manager 142 to stage a track into the cache 140, which may be invoked at block 904 in FIG. 9 when the target track of a read/write request is not in the cache 140. Upon initiating (at block 1000) the operation to stage a track into the cache 140, if (at block 1002) the cache LRU list 400 is full, then the track at the LRU end 404 of the cache LRU list 400 is selected (at block 1004) to demote. If (at block 1006) the demoted cache LRU

list **500** is full, then the cache manager **142** selects (at block **1008**) a demoted track indicated at the LRU end **504** of the demoted cache LRU list **500** to demote. The selected demoted track is removed (at block **1010**) from the LRU end **504**. The cache manager **142** adjusts (at block **1012**) a demoted cache control block **600_j**, whose pointer to next LRU entry **608** points to the selected demoted track in the demoted cache LRU list **500** to indicate that the pointer **608** is null, because now that entry is at the LRU end **504** when the selected demoted track is removed from the demoted cache LRU list **500**.

The cache manager **142** determines (at block **1014**) an entry **700_i** in the demoted cache control block directory **700** having the demoted cache control block **600_s** of the selected demoted track. In one embodiment, the entry **700_i** may be determined by applying a hash function to the selected demoted track identifier. The hash function may map any of the track identifiers in the storage **104** to one of the entries **700_i** in the demoted cache control block directory **700**. The cache manager **142** then needs to adjust the pointer **610** that points to the demoted cache control block **600_s** of the selected demoted track. For this, the cache manager **142** adjusts (at block **1016**) a demoted cache control block **600_j** in the determined entry **700_i** that points to the demoted cache control block **600_s** of the selected demoted track to point to the demoted cache control block **600_k** pointed to by the pointer **610** of the demoted cache control block **600_s** of the selected demoted track. The demoted cache control block **600_s** for the selected demoted track is deleted (at block **1018**) and indication of the deleted demoted cache control block **600_s** is removed (at block **1020**) from the entry **700_i** in the demoted cache control block directory **700**.

From block **1020** or if (at block **1006**) the demoted cache LRU list **500** is not full, control proceeds (at block **1022**) to block **1030** in FIG. **10b** to add the demoted track from the cache **140** to the demoted cache LRU list **500**. Upon initiating (at block **1030**) the operation to add the demoted track to the demoted cache LRU list **500**, the cache manager **142** indicates (at block **1032**) the demoted track at the MRU end **502** of the demoted cache LRU list **500**. The cache manager **142** determines (at block **1034**) the cache control block **300_{DT}** for the demoted track from the cache **140** and the track format code **314** for the demoted track if one is included. The cache manager **142** generates (at block **1036**) a demoted track cache control block **600_{DT}** for the track being demoted indicating a track identifier **602** of the demoted track and the determined track format code **314** for the demoted track to save the track format code in field **604** of the demoted track cache control block **600_{DT}**. If there was no valid track format code **314** in the cache control block **300_{DT}**, then a null value may be indicated in the field **604** indicating there is no valid track format code for the demoted track.

The cache manager **142** determines (at block **1038**) the entry **700_i** in the demoted cache control block directory **700** that will be used to indicate the demoted cache control block **600_{DT}** of the demoted track. The next pointer **610** of the last demoted cache control block **600_i** indicated in the determined entry **700_i** is adjusted (at block **1040**) to point to the demoted cache control block **600_{DT}** for the track being demoted. The demoted track control block **600_{DT}** for the demoted track is indicated in the determined entry **700_i** in the demoted cache control block directory **700**, e.g., at the end of the entry **700_i**. After adding the demoted track to the demoted LRU list **500** and the demoted cache control block **600_{DT}** to the demoted cache control block directory **700**, control proceeds (at block **1044**) to block **1060** in FIG. **10c**

to stage the target track to the cache **140**. Control also proceeds to block **1060** in FIG. **10c** to stage the track if (at block **1002** in FIG. **10a**) the cache LRU list **400** is not full, so that a track does not need to be demoted from the cache **140** and added to the demoted cache LRU list **500**.

Upon initiating (at block **1060**) the operation to stage the track to the cache **140**, the cache manager **142** stages (at block **1062**) the target track into the cache **140** from the storage **104**. The staging of the actual track data from the storage **104** may have been initiated earlier before or during the operations of FIGS. **10a**, **10b** to manage the demote track LRU list **500** and demote cache control block directory **700**. The target track staged into the cache **140** is indicated (at block **1064**) at the MRU end **402** of the cache LRU list **400**. A cache control block **300_{ST}** is generated (at block **1066**) for the staged track. Control then proceeds to block **1068** to determine if there is a demoted cache control block **600_{ST}** for the staged track that has a track format code **604** (or other track format information) that can be included in the cache control block **300_{ST}** created for the staged track. A determination is made (at block **1068**) of the entry **700_i** in the demoted cache control block directory **700** that could have a demoted cache control block **600_{ST}** for the staged track, which entry **700_i** may be determined by applying a hash function to the track identifier of the staged track. The first demoted cache control block **600_{SEL}** in the determined entry **700_i** is selected (at block **1070**). If (at block **1072**) the track identifier **602** of the selected demoted cache control block **600_{SEL}** matches the track identifier of the staged track, then the track format code **604** in the selected demoted cache control block **600_{SEL}** is included (at block **1074**) in the cache control block **300_{ST}** for the staged track. The cache manager **142** may then perform (at block **1076**) the operations at blocks **1010**, **1016**, **1018**, and **1020** in FIG. **10a** to remove demoted track information for the staged track, including removing the demoted cache control block **600_{ST}** for the staged track, removing the staged track from the demoted cache LRU list **500**, removing the indication of the demoted cache control block **600_{ST}** from the demoted cache control block directory **700**, and adjusting pointers **606**, **608**, **610** in other demoted cache control blocks **600_i** that pointed to the demoted track or demoted cache control block **600_{ST}** for the staged track because the staged track is no longer demoted but active in cache **140**.

If (at block **1072**) the selected demoted cache control block **600_{SEL}** is not for the staged track and if (at block **1078**) there is a next demoted cache control block **600_i** in the entry **700_i**, which may be indicated in the pointer **610**, then that next demoted cache control block is selected (at block **1080**) and control returns to block **1070** to determine whether this next demoted cache control block **600_i** is for the staged track. If (at block **1078**) there are no further next demoted cache control blocks in the determined entry **700_i** to consider, then the track format code **202** from a demoted track information cannot be used and the cache manager **142** reads (at block **1082**) the metadata for the track from the storage **104** to determine the track format. From block **1076** after using the track format code **604** from the demoted cache control block for the staged track or after reading (at block **1082**) the metadata for the staged track, control returns to block **904** in FIG. **9** with staging complete to perform the read/write operation with respect to the staged track.

With the embodiments of FIGS. **10a**, **10b**, and **10c**, the track format code for a track demoted from cache can be saved and later used when the demoted track is staged back into cache. This allows the track metadata format to be quickly determined for the demoted track staged back into

11

cache without having to read the metadata for the track from storage. The computer performance for cache operations, particularly staging, are substantially improved and latency reduced by determining the track metadata format and layout of a track being staged into cache without having to read the metadata for the track.

FIG. 11 illustrates an embodiment of operations performed by the cache manager 142 when closing the track metadata for a track in the cache 140, which involves destaging the track metadata to the storage 104 if changed. Upon closing (at block 1100) the track metadata for a track in the cache 140, the cache manager 140 processes (at block 1102) the track metadata to determine a track format or a layout of data in the track. If (at block 1104) the track format table 200 does not have a track format 204 matching the determined track format from the track metadata, which may happen if the determined track format is irregular, then the track format code valid flag 316 is set (at block 1106) to invalid and the invalid reason 318 is set to indicate that the track format is not supported. In such situation, read/write requests to the track having an irregular format are only processed when received through the second channel via network host adaptor 134.

If (at block 1104) the track format table has a track format 204 matching the determined track format from the track metadata, then the cache manager 142 determines the track format code 202 for the determined track format 204 in the track format table 200 and includes the track format code 202 in the field 314 in the cache control block 300_i. The track format code valid flag 316 is set (at block 1116) to valid. From block 1108 or 1116, control proceeds to block 1118 to destage the track metadata from the memory 138 if modified or discard if not modified.

With the operations of FIG. 11, the track format information may be indicated in the cache control block 300_i with a track format code 202 having a limited number of bits to index a track format descriptor 204 or track layout in a track format table 200, where the track metadata itself would not fit into the cache control block 300_i. For future read/write accesses, if a valid track format code 314 is provided, then the cache manager 142 may use that code 314 to obtain with low latency the track format metadata 204 from the track format table 200 without having to read the track metadata from the storage 104 and process to determine the track format.

Invalidating a Metadata Track

A metadata track stored in the storage 104 and the cache 140 when staged into the cache 140 includes metadata for a range of tracks in the storage 104. A metadata track may maintain a flag for each track having metadata in the metadata track indicating whether the metadata for the track in the metadata track is valid or invalid. Metadata tracks may be invalidated when there is a release space operation performed to free tracks or remove them from a volume or when there is a change to the track layout and format, resulting in the current track metadata being invalidated.

FIG. 12 illustrates an embodiment of operations performed by the cache manager 142 to invalidate track format information, such as track format codes 604 saved for tracks in cache 140 or demoted from the cache 140. Upon initiating (at block 1200) an operation to invalidate a metadata track, a determination is made (at block 1202) of tracks having metadata in the metadata track to invalidate. If (at block 1204) the cache control block directory 300 includes cache control blocks 300_i for any of the determined tracks having

12

metadata in the metadata track to invalidate, then the track format code valid flag 316 in each of the determined cache control blocks is set (at block 1206) to indicate the track format code 314 is invalid. If (at block 1204) the cache control block directory 300 does not include cache control blocks for any of the determined tracks or after setting (at block 1206) the valid flag 316 to invalid, the cache manager 142 sets (at block 1208) a track i to the first track in the range of tracks having metadata in the metadata track to invalidate. In certain embodiments, the metadata track maintains metadata information for a range of consecutive tracks configured in a volume in the storage 104.

The cache manager 142 determines (at block 1210) (e.g., from hash function of the track i) an entry 700_j in the demoted cache control block directory 700 that would have a demoted cache control block 600_i for track i. If (at block 1212) the determined entry 700_i indicates a demoted cache control block 600_i for track i, i.e., having a track ID 602 matching track i, then the cache manager 142 removes (at block 1214) the demoted cache control block 600_i for track i and removes indication of the removed demoted cache control block 600_i from the determined entry 700_i in the demoted cache control block directory 700. The cache manager 142 further removes (at block 1216) the entry for track i in the demoted cache LRU list 500. If (at block 1212) the determined entry 700_i for track i does not have a demoted cache control block 600_i for track i or from block 1216, if (at block 1218) there is a next track in the range of tracks having metadata in the metadata track to invalidate, then track i is incremented (at block 1220) to a next track having metadata in the metadata track to invalidate and control proceeds to block 1210 to consider the next track i having metadata in the metadata track to invalidate. If (at block 1218) there are no further tracks in the range to consider, then the metadata track is invalidated (at block 1222).

With the operations of FIG. 12, the metadata track is invalidated only after invalidating track format information, e.g., the track format code 314, in the cache control blocks 300_i for the tracks having metadata in the metadata track to invalidate or after removing demoted cache control blocks having the track form information, e.g., track format code 604, for tracks having metadata in the metadata track to invalidate. In this way, all instances of cache information in cache control blocks 300_i and demoted cache control blocks 600_i that would have the saved track format information based on track metadata are removed before invalidating the track metadata. This ensures that the cache manager 142 will not rely on invalid track format information when the metadata track is invalidated because the track format information is removed from cache information, e.g., control blocks 300_i and 600_i before the track metadata is invalidated. The track metadata may be invalidated by setting flags in the track metadata for each track having metadata in the metadata track to invalidate. After the metadata track is invalidated, the track metadata would have to be rebuilt by accessing the track and processing the track to determine the track metadata and update the metadata track with the rebuilt track metadata.

FIG. 13 illustrates an embodiment of operations performed by the cache manager 142 to call functions to invalidate cache control blocks 300_i and demoted cache control blocks 600_i that have saved track format information for tracks in a metadata track to invalidate. Upon initiating (at block 1300) an operation to invalidate a metadata track, the cache manager 142 calls (at block 1302) a cache control block function with the range of tracks in the metadata track to invalidate to perform operations at blocks to perform

operations at blocks 1202-1206 in FIG. 12 to invalidate the track format information (track format codes 314) in cache control blocks 300, for tracks having metadata in the metadata track to invalidate. The cache manager 142 further calls (at block 1304) a demoted cache control block function with the range of tracks in the metadata track to perform operations at blocks 1208-1220 in FIG. 12 to remove tracks from the demoted track list and demoted cache control blocks for the tracks having metadata in the metadata track to invalidate. After invalidating all the track format information for the tracks having track metadata in the metadata track to invalidate, the metadata track is then invalidated (at block 1306), in the cache 140 or the storage 104.

The present invention may be implemented as a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like, and conventional procedural programming

languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be

noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

The computational components of FIG. 1, including the host 100 and storage system 102 may be implemented in one or more computer systems, such as the computer system 1402 shown in FIG. 14. Computer system/server 1402 may be described in the general context of computer system executable instructions, such as program modules, being executed by a computer system. Generally, program modules may include routines, programs, objects, components, logic, data structures, and so on that perform particular tasks or implement particular abstract data types. Computer system/server 1402 may be practiced in distributed cloud computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed cloud computing environment, program modules may be located in both local and remote computer system storage media including memory storage devices.

As shown in FIG. 14, the computer system/server 1402 is shown in the form of a general-purpose computing device. The components of computer system/server 1402 may include, but are not limited to, one or more processors or processing units 1404, a system memory 1406, and a bus 1408 that couples various system components including system memory 1406 to processor 1404. Bus 1408 represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnects (PCI) bus.

Computer system/server 1402 typically includes a variety of computer system readable media. Such media may be any available media that is accessible by computer system/server 1402, and it includes both volatile and non-volatile media, removable and non-removable media.

System memory 1406 can include computer system readable media in the form of volatile memory, such as random access memory (RAM) 1410 and/or cache memory 1412. Computer system/server 1402 may further include other removable/non-removable, volatile/non-volatile computer system storage media. By way of example only, storage system 1413 can be provided for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a "hard drive"). Although not shown, a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a "floppy disk"), and an optical disk drive for reading from or writing to a removable, non-volatile optical disk such as a CD-ROM, DVD-ROM or other optical media can be provided. In such instances, each can be connected to bus 1408 by one or more data media interfaces. As will be further depicted and described below, memory 1406 may include at least one program product having a set (e.g., at least one) of program modules that are configured to carry out the functions of embodiments of the invention.

Program/utility 1414, having a set (at least one) of program modules 1416, may be stored in memory 806 by way of example, and not limitation, as well as an operating

system, one or more application programs, other program modules, and program data. Each of the operating system, one or more application programs, other program modules, and program data or some combination thereof, may include an implementation of a networking environment. The components of the computer 1402 may be implemented as program modules 1416 which generally carry out the functions and/or methodologies of embodiments of the invention as described herein. The systems of FIG. 1 may be implemented in one or more computer systems 1402, where if they are implemented in multiple computer systems 1402, then the computer systems may communicate over a network.

Computer system/server 1402 may also communicate with one or more external devices 1418 such as a keyboard, a pointing device, a display 1420, etc.; one or more devices that enable a user to interact with computer system/server 1402; and/or any devices (e.g., network card, modem, etc.) that enable computer system/server 1402 to communicate with one or more other computing devices. Such communication can occur via Input/Output (I/O) interfaces 1422. Still yet, computer system/server 1402 can communicate with one or more networks such as a local area network (LAN), a general wide area network (WAN), and/or a public network (e.g., the Internet) via network adapter 1424. As depicted, network adapter 1424 communicates with the other components of computer system/server 1402 via bus 1408. It should be understood that although not shown, other hardware and/or software components could be used in conjunction with computer system/server 1402. Examples, include, but are not limited to: microcode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

The terms "an embodiment", "embodiment", "embodiments", "the embodiment", "the embodiments", "one or more embodiments", "some embodiments", and "one embodiment" mean "one or more (but not all) embodiments of the present invention(s)" unless expressly specified otherwise.

The terms "including", "comprising", "having" and variations thereof mean "including but not limited to", unless expressly specified otherwise.

The enumerated listing of items does not imply that any or all of the items are mutually exclusive, unless expressly specified otherwise.

The terms "a", "an" and "the" mean "one or more", unless expressly specified otherwise.

Devices that are in communication with each other need not be in continuous communication with each other, unless expressly specified otherwise. In addition, devices that are in communication with each other may communicate directly or indirectly through one or more intermediaries.

A description of an embodiment with several components in communication with each other does not imply that all such components are required. On the contrary a variety of optional components are described to illustrate the wide variety of possible embodiments of the present invention.

When a single device or article is described herein, it will be readily apparent that more than one device/article (whether or not they cooperate) may be used in place of a single device/article. Similarly, where more than one device or article is described herein (whether or not they cooperate), it will be readily apparent that a single device/article may be used in place of the more than one device or article or a different number of devices/articles may be used instead of the shown number of devices or programs. The functionality and/or the features of a device may be alternatively embod-

ied by one or more other devices which are not explicitly described as having such functionality/features. Thus, other embodiments of the present invention need not include the device itself.

The foregoing description of various embodiments of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be limited not by this detailed description, but rather by the claims appended hereto. The above specification, examples and data provide a complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the spirit and scope of the invention, the invention resides in the claims herein after appended.

What is claimed is:

1. A computer program product for managing metadata on metadata tracks in a storage, wherein the metadata tracks provide metadata on tracks cached in a cache in a portion of a memory, the computer program product comprising a computer readable storage medium having computer readable program code embodied therein that is executable to perform operations, the operations comprising:

determining tracks having metadata in a metadata track to invalidate; and

indicating track format codes, in cache control blocks for tracks whose metadata is in the metadata track to invalidate, as invalid, wherein valid track format codes are used to determine corresponding track format metadata having track format layout information used to access tracks, wherein the track format codes identify format metadata, wherein there are multiple instances of the track format codes maintained separate from the metadata tracks.

2. The computer program product of claim 1, wherein the operations further comprise:

receiving a read/write request to a target track on a first channel;

failing the read/write request on the first channel in response to a track format code, in a cache control block for the target track, indicated as invalid; and

processing the read/write request on the first channel using track format layout information in track format metadata corresponding to the track format code, in the cache control block for the target track, in response to the track format code for the target track indicated as valid.

3. The computer program product of claim 2, wherein the operations further comprise:

receiving the read/write request to the target track on a second channel in response to failing the read/write request on the first channel; and

processing the read/write request on the second channel using track format layout information read from metadata for the target track in the storage in response to the track format code, in the cache control block for the target track, indicated as invalid.

4. The computer program product of claim 1, wherein the operations further comprise:

maintaining a track format table in the memory associating the track format codes with different track format metadata, wherein each instance of the track format metadata in the track format table indicates a layout of data in a track.

5. The computer program product of claim 4, wherein each track format code of the track format codes comprises an index value of an entry in the track format table having the track format metadata associated with the track format code.

6. The computer program product of claim 1, wherein each cache control block of the cache control blocks includes a valid flag indicating whether the cache control block includes a valid track format code, wherein the valid flag is set to valid when adding a track format code of the track format codes to the cache control block, wherein the indicating a track format code in a cache control block as invalid comprises:

setting a valid flag to invalid in the cache control block for a track whose metadata is in the metadata track to invalidate.

7. The computer program product of claim 1, wherein the metadata track to invalidate has a valid flag for each track having metadata in the metadata track to invalidate indicating whether the metadata for the track is valid or invalid, wherein the operations further comprise: setting the valid flag for a track in the metadata track to invalidate to invalid in response to indicating a track format code of the track format codes in a cache control block for the track as invalid; and invalidating the metadata track to invalidate after indicating all the track format codes as invalid.

8. A system for managing metadata on metadata tracks in a storage, comprising:

a processor;

a memory including a cache in a portion of the memory to cache tracks from the storage, wherein the metadata tracks provide metadata on tracks cached in the cache; a computer readable storage medium having computer readable program code embodied therein that is executable to perform operations, the operations comprising: determining tracks having metadata in a metadata track to invalidate; and

indicating track format codes, in cache control blocks for tracks whose metadata is in the metadata track to invalidate, as invalid, wherein valid track format codes are used to determine corresponding track format metadata having track format layout information used to access tracks, wherein the track format codes identify format metadata, wherein there are multiple instances of the track format codes maintained separate from the metadata tracks.

9. The system of claim 8, wherein the operations further comprise:

receiving a read/write request to a target track on a first channel;

failing the read/write request on the first channel in response to a track format code, in a cache control block for the target track, indicated as invalid; and

processing the read/write request on the first channel using track format layout information in track format metadata corresponding to the track format code, in the cache control block for the target track, in response to the track format code for the target track indicated as valid.

10. The system of claim 9, wherein the operations further comprise:

receiving the read/write request to the target track on a second channel in response to failing the read/write request on the first channel; and

processing the read/write request on the second channel using track format layout information read from metadata for the target track in the storage in response to the

19

track format code, in the cache control block for the target track, indicated as invalid.

11. The system of claim 8, wherein the operations further comprise:

maintaining a track format table in the memory associating the track format codes with different track format metadata, wherein each instance of the track format metadata in the track format table indicates a layout of data in a track.

12. The system of claim 11, wherein each track format code of the track format codes comprises an index value of an entry in the track format table having the track format metadata associated with the track format code.

13. The system of claim 8, wherein each cache control block of the cache control blocks includes a valid flag indicating whether the cache control block includes a valid track format code, wherein the valid flag is set to valid when adding a track format code of the track format codes to the cache control block, wherein the indicating a track format code in a cache control block as invalid comprises:

setting a valid flag to invalid in the cache control block for a track whose metadata is in the metadata track to invalidate.

14. The system of claim 8, wherein the metadata track to invalidate has a valid flag for each track having metadata in the metadata track to invalidate indicating whether the metadata for the track is valid or invalid, wherein the operations further comprise: setting the valid flag for a track in the metadata track to invalidate to invalid in response to indicating a track format code of the track format codes in a cache control block for the track as invalid; and invalidating the metadata track to invalidate after indicating all the track format codes as invalid.

15. A method for managing metadata on metadata tracks in a storage, wherein the metadata tracks provide metadata on tracks cached in a cache in a memory, comprising:

determining tracks having metadata in a metadata track to invalidate; and

indicating track format codes, in cache control blocks for tracks whose metadata is in the metadata track to invalidate, as invalid, wherein valid track format codes are used to determine corresponding track format metadata having track format layout information used to access tracks, wherein the track format codes identify format metadata, wherein there are multiple instances of the track format codes maintained separate from the metadata tracks.

16. The method of claim 15, further comprising: receiving a read/write request to a target track on a first channel;

20

failing the read/write request on the first channel in response to a track format code, in a cache control block for the target track, indicated as invalid; and processing the read/write request on the first channel using track format layout information in track format metadata corresponding to the track format code, in the cache control block for the target track, in response to the track format code for the target track indicated as valid.

17. The method of claim 16, further comprising: receiving the read/write request to the target track on a second channel in response to failing the read/write request on the first channel; and

processing the read/write request on the second channel using track format layout information read from metadata for the target track in the storage in response to the track format code, in the cache control block for the target track, indicated as invalid.

18. The method of claim 15, further comprising: maintaining a track format table in the memory associating the track format codes with different track format metadata, wherein each instance of the track format metadata in the track format table indicates a layout of data in a track.

19. The method of claim 18, wherein each track format code of the track format codes comprises an index value of an entry in the track format table having the track format metadata associated with the track format code.

20. The method of claim 15, wherein each cache control block of the cache control blocks includes a valid flag indicating whether the cache control block includes a valid track format code, wherein the valid flag is set to valid when adding a track format code of the track format codes to the cache control block, wherein the indicating a track format code in a cache control block as invalid comprises:

setting a valid flag to invalid in the cache control block for a track whose metadata is in the metadata track to invalidate.

21. The method of claim 15, wherein the metadata track to invalidate has a valid flag for each track having metadata in the metadata track to invalidate indicating whether the metadata for the track is valid or invalid, further comprising: setting the valid flag for a track in the metadata track to invalidate to invalid in response to indicating a track format code of the track format codes in a cache control block for the track as invalid; and invalidating the metadata track to invalidate after indicating all the track format codes as invalid.

* * * * *